Addendum for Dell Networking OS 9.3(0.0)

Notes, Cautions, and Warnings

NOTE: A NOTE indicates important information that helps you make better use of your computer.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

© 2013 Dell Inc. All Rights Reserved.

Trademarks used in this text: Dell™, the Dell logo, Dell Boomi™, Dell Precision™ , OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™, Venue™ and Vostro™ are trademarks of Dell Inc. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, vMotion®, vCenter®, vCenter SRM™ and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation.

2014 - 02

Rev. A00

Contents

1 About this Document....23
Audience....23
Conventions....23
Related Documents....24

2 802.1X on the MXL 10/40GbE Switch....25

3 ACL VLAN Groups and Content Addressable Memory (CAM)....27
Optimizing CAM Utilization During the Attachment of ACLs to VLANs....27
Guidelines for Configuring ACL VLAN groups....28
Configuring ACL VLAN Groups and Configuring FP Blocks for VLAN Parameters....29
Configuring ACL VLAN Groups....29
Configuring FP Blocks for VLAN Parameters....30
Viewing CAM Usage....31
Allocating FP Blocks for VLAN Processes....32
member vlan....33
ip access-group....34
show acl-vlan-group....34
show cam-acl-vlan....36
cam-acl-vlan....37
show cam-usage....38
show running config acl-vlan-group....41
acl-vlan-group....41
show acl-vlan-group detail....42
description (ACL VLAN Group)....43

4 Access Control Lists....45
Logging of ACL Processes....45
Guidelines for Configuring ACL Logging....46
Configuring ACL Logging....47
deny (for Standard IP ACLs)....48
deny (for Extended IP ACLs)....49
seq (for Standard IPv4 ACLs)....50
deny tcp (for Extended IP ACLs)....51
deny udp (for Extended IP ACLs)....52
deny arp (for Extended MAC ACLs)....53
deny icmp (for Extended IP ACLs)....54
deny ether-type (for Extended MAC ACLs)....56

deny (for Standard MAC ACLs)....57
deny (for Extended MAC ACLs)....58
permit arp (for Extended MAC ACLs)....59
permit ether-type (for Extended MAC ACLs)....60
permit icmp (for Extended IP ACLs)....61
permit udp (for Extended IP ACLs)....62
permit (for Extended IP ACLs)....63
permit (for Standard MAC ACLs)....65
seq (for Standard MAC ACLs)....66
permit tcp (for Extended IP ACLs)....67
seq arp (for Extended MAC ACLs)....68
seq ether-type (for Extended MAC ACLs)....69
seq (for IP ACLs)....70
seq (for IPv6 ACLs)....71
permit udp (for IPv6 ACLs)....72
permit tcp (for IPv6 ACLs)....73
permit icmp (for IPv6 ACLs)....75
permit (for IPv6 ACLs)....76
deny udp (for IPv6 ACLs)....77
deny tcp (for IPv6 ACLs)....78
deny icmp (for Extended IPv6 ACLs)....79
deny (for IPv6 ACLs)....80
Flow-Based Monitoring Support for ACLs....81
Behavior of Flow-Based Monitoring....82
Enabling Flow-Based Monitoring....84

5 **Bare Metal Provisioning (BMP)....85**
Support for BMP on the S6000 Switch....85
Enhanced Behavior of the stop bmp Command....85
Removal of the Deprecated User-Defined String Parameter With reload-type Command....85
Inclusion of Service Tag Information in the Option 60 String....85
Replacement of stop jump-start Command With the stop bmp Command....86

6 **Data Center Bridging (DCB)....87**
Configuring DCB Maps and its Attributes....87
DCB Map: Configuration Procedure....87
Important Points to Remember....88
Applying a DCB Map on a Port....88
Configuring PFC without a DCB Map....89
Configuring Lossless Queues....89
Data Center Bridging: Default Configuration....90
Configuring PFC and ETS in a DCB Map....91

PFC Configuration Notes 91
PFC Prerequisites and Restrictions 92
ETS Configuration Notes 92
ETS Prerequisites and Restrictions 93
dcb-map 94
S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator 94
priority-pgid 95
S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator 95
pfc mode on 96
priority-group bandwidth pfc 97
S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator 97
dcb-map stack-unit all stack-ports all 98
S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator 98
show qos dcb-map 99
S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator 99
Priority-Based Flow Control Using Dynamic Buffer Method 100
Pause and Resume of Traffic 100
Buffer Sizes for Lossless or PFC Packets 100
Interworking of DCB Map With DCB Buffer Threshold Settings 101
Configuring the Dynamic Buffer Method 102
Applying a DCB Map in a Switch Stack 103
dcb pfc-shared-buffer-size 103
S6000 S4810 S4820T MXL 103
dcb-buffer-threshold 104
S6000 S4810 S4820T MXL 104
priority 105
S6000 S4810 S4820T MXL 105
qos-policy-buffer 106
S6000 S4810 S4820T MXL 106
dcb-policy buffer-threshold (Interface Configuration) 108
S6000 S4810 S4820T MXL 108
dcb-policy dcb-buffer-threshold (Global Configuration) 109
S4810 S4820T MXL 109
show qos dcb-buffer-threshold 109
show hardware stack-unit buffer-stats-snapshot (With Polling and History) 110
dcb pfc-total-buffer-size 117
S6000 117
show running-config dcb-buffer-threshold 117
dcb pfc-queues 119

7 Egress Interface Selection (EIS) for HTTP and IGMP Applications 121

Protocol Separation 121

Enabling and Disabling Management Egress Interface Selection....122
Handling of Management Route Configuration....123
Handling of Switch-Initiated Traffic....124
Handling of Switch-Destined Traffic....125
Handling of Transit Traffic (Traffic Separation)....125
Mapping of Management Applications and Traffic Type....126
Behavior of Various Applications for Switch-Initiated Traffic....127
Behavior of Various Applications for Switch-Destined Traffic....128
Interworking of EIS With Various Applications....128
application (for HTTP and ICMP)....129
Z9000 S4810 S4820T....129

8 Flex Hash and Optimized Boot-Up....131
Flex Hash Capability Overview....131
load-balance ingress-port enable....132
load-balance flexhash....132
Configuring the Flex Hash Mechanism....134
Configuring Fast Boot and LACP Fast Switchover....135
reload-type fastboot....135
S6000....135
lacp fast-switchover....136
S6000....136
Optimizing the Boot Time....136
Booting Process When Optimized Boot Time Mechanism is Enabled....137
Guidelines for Configuring Optimized Booting Mechanism....137
Interoperation of Applications with Fast Boot and System States....138
LACP and IPv4 Routing....139
LACP and IPv6 Routing....139
BGP Graceful Restart....140
Cold Boot Caused by Power Cycling the System....140
Unexpected Reload of the System....140
Software Upgrade....140
LACP Fast Switchover....141
Changes to BGP Multipath....141
Minimized Connection Setup Time....141
Faster Local Route Aadvertisements....141
Delayed Installation of ECMP Routes Into BGP....142
Changes for BGP Graceful Restart Processes....142
Operation of LACP....142
Operation of FIB....143
RDMA Over Converged Ethernet (RoCE) Overview....143
Preserving 802.1Q VLAN Tag Value for Lite Subinterfaces....144

encapsulation dot1q....145

9 Interfaces....147
Enabling the Management Address TLV on All Interfaces of an Aggregator....147
Enhanced Validation of Interface Ranges....147

10 IPv4 Routing....149
IPv4 Path MTU Discovery Overview....149
Using the Configured Source IP Address in ICMP Messages....150
Configuring the ICMP Source Interface....150
Working of the Traceroute Utility....150
ip icmp source-interface....151
ipv6 icmp source-interface....152
Configuring the Duration to Establish a TCP Connection....154
ip tcp initial-time....154
show ip tcp initial-time....155

11 Link Aggregation Groups (LAGs)....157
Configuring the Minimum Number of Links to be Up for Uplink LAGs to be Active....157
Optimizing Traffic Disruption Over LAG Interfaces On IOA Switches in VLT Mode....158
Preserving LAG and Port Channel Settings in Nonvolatile Storage....158
Enabling the Verification of Member Links Utilization in a LAG Bundle....159
Monitoring the Member Links of a LAG Bundle....159
show link-bundle-distribution port-channel....160
Setting Up a Threshold for Utilization of High-Gigabit Port Channels....161
Guidelines for Configuring the Mechanism to Monitor High-Gigabit Port Channels....162
Enabling the Verification of Member Links Utilization in a High-Gigabit Port Channel....163
hg-link-bundle-monitor....164
hg-link-bundle-monitor trigger-threshold....165
hg-link-bundle-monitor rate-interval....165
show hg-link-bundle-distribution....166
snmp-server enable traps (for High-Gigabit Port Channel)....167
show hardware stack-unit (for high-Gigabit Ethernet ports)....167
Z9000....168
clear hardware stack-unit (for high-Gigabit Ethernet ports)....169
Z9000....169
Viewing Buffer Utilization and Queue Statistics on High-Gigabit Ethernet Backplane Ports....170

12 Miscellaneous Settings....173
Setting a Threshold for Switching to the SPT....173
ip pim spt-threshold....173
S6000....173

ip route bfd (for S6000)..... 174
S6000..... 174
Configure BFD for Static Routes..... 175
Related Configuration Tasks..... 175
Changing Static Route Session Parameters..... 175
Establishing Sessions for Static Routes..... 176
Disabling BFD for Static Routes..... 176
source (port monitoring for 40-Gigabit Ethernet)..... 177

13 Microsoft Network Load Balancing..... 179
NLB Unicast Mode Scenario..... 179
NLB Multicast Mode Scenario..... 180
Limitations With Enabling NLB on Switches..... 180
Benefits and Working of Microsoft Clustering..... 180
Enable and Disable VLAN Flooding..... 180
Configuring a Switch for NLB..... 181
..... 181
arp (for Multicast MAC Address)..... 181
mac-address-table static (for Multicast MAC Address)..... 182
ip vlan-flooding..... 184

14 Quality of Service (QoS)..... 185
Specifying Policy-Based Rate Shaping in Packets Per Second..... 185
Configuring Policy-Based Rate Shaping..... 186
Configuring Weights and ECN for WRED..... 186
Global Service Pools With WRED and ECN Settings..... 187
Configuring WRED and ECN Attributes..... 188
Classifying Layer 2 Traffic on Layer 3 Interfaces..... 189
Managing Hardware Buffer Statistics..... 190
Enabling Buffer Statistics Tracking..... 191
Classifying Packets Based on a Combination of DSCP Code Points and VLAN IDs..... 191
rate shape..... 192
S6000..... 192
buffer-stats-snapshot..... 194
S6000..... 194
service-class buffer shared-threshold-weight..... 195
S6000Z9000..... 195
wred weight..... 197
S6000Z9000..... 197
service-class wred..... 197
Z9000..... 197
service-pool wred..... 199

S6000Z9000 ... 199
service-class wred ... 200
service-class wred ecn ... 201
Z9000 ... 201
show hardware stack-unit buffer ... 202
show hardware stack-unit buffer-stats-snapshot ... 204
show hardware stack-unit buffer-stats-snapshot (Total Buffer Information) ... 206

15 Management Port Media Converter ... 209
Management Port Media Converter Components ... 209
Working of the Management Port Media Converter ... 210
Online Insertion and Removal (OIR) of the Management Optic ... 212

16 Security for M I/O Aggregator ... 215
aaa authentication enable ... 215
aaa authentication login ... 216
access-class ... 217
Authorization and Privilege Commands ... 218
banner exec ... 218
banner login ... 219
banner motd ... 220
debug radius ... 221
debug tacacs+ ... 221
enable secret ... 221
exec-banner ... 222
ip radius source-interface ... 223
ip tacacs source-interface ... 223
login authentication ... 224
motd-banner ... 225
password-attributes ... 225
privilege level (CONFIGURATION mode) ... 226
privilege level (LINE mode) ... 227
RADIUS Commands ... 227
radius-server deadtime ... 227
radius-server host ... 228
radius-server retransmit ... 229
radius-server timeout ... 229
radius-server key ... 230
show privilege ... 231
Suppressing AAA Accounting for Null Username Sessions ... 231
TACACS+ Commands ... 231
tacacs-server host ... 231

tacacs-server key....232
timeout login response....233
Understanding Banner Settings....233
AAA Authentication....234
Configuration Task List for AAA Authentication....234
RADIUS....236
RADIUS Authentication and Authorization....237
Configuration Task List for RADIUS....238
TACACS+....241
Configuration Task List for TACACS+....241
TACACS+ Remote Authentication and Authorization....242
Command Authorization....244
Protection from TCP Tiny and Overlapping Fragment Attacks....244
Enabling SCP and SSH....244
Using SCP with SSH to Copy a Software Image....245
Secure Shell Authentication....246
Troubleshooting SSH....248
Telnet....249
VTY Line and Access-Class Configuration....249
VTY Line Local Authentication and Authorization....249
VTY Line Remote Authentication and Authorization....250
VTY MAC-SA Filter Support....251

17 Simple Network Management Protocol (SNMP)....253
SNMPv3 Compliance With FIPS....253
snmp-server user (for AES128-CFB Encryption)....254
Z-Series S4810 S4820T S6000 MXL I/O Aggregator....254

18 Stacking....257
Configuring the Uplink Speed of Interfaces as 40 Gigabit Ethernet....257
stack-unit iom-mode uplink-speed....258
show system stack-unit iom-uplink-speed....259
stack-unit priority....260
stack-unit renumber....260

19 Virtual Link Trunking (VLT)....263
Specifying VLT Nodes in a PVLAN....263
Association of VLTi as a Member of a PVLAN....264
MAC Synchronization for VLT Nodes in a PVLAN....264
PVLAN Operations When One VLT Peer is Down....265
PVLAN Operations When a VLT Peer is Restarted....265
Interoperation of VLT Nodes in a PVLAN with ARP Requests....265

Scenarios for VLAN Membership and MAC Synchrnoization With VLT Nodes in PVLAN.....265
Configuring a VLT VLAN or LAG in a PVLAN.....267
Creating a VLT LAG or a VLT VLAN..... 267
Associating the VLT LAG or VLT VLAN in a PVLAN.....268
show vlt private-vlan..... 269
Proxy ARP Capability on VLT Peer Nodes.....270
Working of Proxy ARP for VLT Peer Nodes.....270
VLT Nodes as Rendezvous Points for Multicast Resiliency.....271

20 Documentation Updates.....273
Configuring the Commands Without a Separate User Account for the PMUX Mode of the I/O Aggregator..... 277

21 Data Center Bridging (DCB).....279
advertise dcbx-appln-tlv.....279
advertise dcbx-tlv.....279
bandwidth-percentage..... 280
dcb-enable..... 281
dcb-input..... 282
dcb-output.....282
dcb-policy input..... 283
dcb-policy input stack-unit stack-ports all..... 284
dcb-policy output.....284
dcb-policy output stack-unit stack-ports all.....285
dcb stack-unit all pfc-buffering pfc-port-count pfc-queues..... 286
dcb stack-unit pfc-buffering pfc-port-count pfc-queues..... 287
dcbx port-role.....287
dcbx version.....288
debug dcbx..... 289
description..... 290
ets mode on.....290
fcoe priority-bits.....291
iscsi priority-bits..... 291
pfc link-delay..... 292
pfc mode on.....292
pfc no-drop queues.....293
pfc priority.....294
priority-group..... 294
priority-group qos-policy.....295
priority-list.....296
qos-policy-output ets.....297
scheduler.....297

set-pgid....298
show dcb....299
show interface dcbx detail....299
show interface ets....302
show interface pfc....305
show interface pfc statistics....308
show qos dcb-input....309
show qos dcb-output....309
show qos priority-groups....310
show stack-unit stack-ports ets details....310
show stack-unit stack-ports pfc details....311

22 FIP Snooping....313
clear fip-snooping database interface vlan....313
clear fip-snooping statistics....313
feature fip-snooping....314
fip-snooping enable....314
fip-snooping fc-map....315
fip-snooping port-mode fcf....315

23 High Availability (HA)....317
redundancy force-failover....317
Z9000 S4810 S4820T....317
show redundancy....318
Z9000 S4810 S4820T....318

24 iSCSI Optimization....323
advertise dcbx-app-tlv....323
iscsi aging time....323
iscsi cos....324
iscsi enable....325
iscsi priority-bits....325
iscsi profile-compellant....325
iscsi target port....326
iSCSI Optimization Prerequisites....326
Configuring iSCSI Optimization....327

25 Interfaces....331
Basic Interface Commands....331
clear counters....331
description....332
flowcontrol....333

interface....335
interface ManagementEthernet....336
interface range....337
interface vlan....339
keepalive....340
mtu....340
negotiation auto....341
portmode hybrid....343
stack-unit portmode....345
Port Channel Commands....346
channel-member....346
interface port-channel....348
minimum-links....349

26 Internet Group Management Protocol (IGMP)....351
IGMP Commands....351
Important Points to Remember....351
ip igmp group-join-limit....351
ip igmp last-member-query-interval....352
ip igmp querier-timeout....353
ip igmp query-interval....354
ip igmp query-max-resp-time....354
ip igmp version....355
IGMP Snooping Commands....356
Important Points to Remember for IGMP Snooping....356
Important Points to Remember for IGMP Querier....356
ip igmp snooping enable....357
ip igmp snooping fast-leave....358
ip igmp snooping last-member-query-interval....359
ip igmp snooping mrouter....359
ip igmp snooping querier....361

27 Layer 2....363
MAC Addressing Commands....363
mac-address-table aging-time....363
mac-address-table static....364
mac-address-table station-move refresh-arp....364

28 Link Aggregation Control Protocol (LACP)....367
lacp long-timeout....367
lacp port-priority....367
port-channel mode....368

port-channel-protocol lacp....369
Configuration Tasks for Port Channel Interfaces....369
Creating a Port Channel....370
Adding a Physical Interface to a Port Channel....370
Reassigning an Interface to a New Port Channel....372
Configuring the Minimum Oper Up Links in a Port Channel....373
Adding or Removing a Port Channel from a VLAN....373
Configuring VLAN Tags for Member Interfaces....374
Deleting or Disabling a Port Channel....374

29 Link Layer Discovery Protocol (LLDP)....375
advertise dot1-tlv....375
advertise dot3-tlv....376
advertise management-tlv....376
clear lldp counters....377
clear lldp neighbors....377
debug lldp interface....378
disable....379
hello....379
mode....380
multiplier....380
Configure LLDP....381
Related Configuration Tasks....381
Important Points to Remember....381
LLDP Compatibility....381
CONFIGURATION versus INTERFACE Configurations....381
Enabling LLDP....382
Disabling and Undoing LLDP....382
Enabling LLDP on Management Ports....383
Disabling and Undoing LLDP on Management Ports....383
Advertising TLVs....383
Viewing the LLDP Configuration....385
Viewing Information Advertised by Adjacent LLDP Agents....385
Configuring LLDPDU Intervals....386
Configuring Transmit and Receive Mode....387
Configuring a Time to Live....388

30 Quality of Service (QoS)....389
Per-Port QoS Commands....389
dot1p-priority....389
rate shape....390
service-class dynamic dot1p....390

service-class dot1p-mapping....392
Z9000 S4810 S4820T....392
service-class bandwidth-percentage....392
Policy-Based QoS Commands....393
bandwidth-percentage....393
clear qos statistics....394
description....395
policy-aggregate....395
policy-map-output....396
qos-policy-output....397
rate police....397
rate shape....398
service-policy output....399
service-queue....399
set....400
show qos policy-map....401
show qos policy-map-output....402
show qos qos-policy-output....402
show qos statistics....403
show qos wred-profile....404
wred....405
wred-profile....406

31 reload-type....407
Z9000 S4810 S4820TS6000....407

32 Simple Network Management Protocol (SNMP) and Syslog....411
SNMP Commands....411
Important Points to Remember....411
snmp-server enable traps....411
snmp-server host....413
Syslog Commands....416
clear logging....416
logging....417
logging buffered....418
logging console....419
logging monitor....420
logging source-interface....421
show logging....422
show logging driverlog stack-unit....424
terminal monitor....424

33 Storm Control 427
Important Points to Remember 427
show storm-control unknown-unicast 427
Z-Series S4810 S4820TS6000 427
storm-control broadcast (Configuration) 428
Z-Series S4810 S4820TS6000 428
storm-control multicast (Configuration) 429
Z-SeriesS4810 S4820TS6000 429
storm-control broadcast (Interface) 430
Z-Series S4810 S4820TS6000 430

34 Uplink Failure Detection (UFD) 433
clear ufd-disable 433
S4810 S4820T 433
debug uplink-state-group 434
S4810 S4820T 434
description 435
S4810 S4820T 435
downstream 436
S4810 S4820T 436
downstream auto-recover 437
S4810 S4820T 437
downstream disable links 438
S4810 S4820T 438
enable 439
S4810 S4820T 439
show running-config uplink-state-group 439
S4810 S4820T 439
show uplink-state-group 440
S4810 S4820T 440
uplink-state-group 442
S4810 S4820T 442
upstream 443
S4810 S4820T 443

35 Virtual Link Trunking (VLT) 445
back-up destination 445
Z9000 S4810 S4820T 445
clear vlt statistics 446
Z9000 S4810 S4820T 446
delay-restore 447

Z-Series S4810 S4820T ... 447
lacp ungroup member-independent ... 448
Z-Series S4810 S4820T ... 448
peer-link port-channel ... 449
Z-Series S4810 S4820T ... 449
primary-priority ... 450
S4810 S4820T ... 450
show vlt mismatch ... 451
Z9000 S4810 S4820TS6000 ... 451
system-mac ... 451
Z-Series S4810 S4820T ... 451
unit-id ... 452
Z-Series S4810S4820T ... 452
vlt domain ... 453
Z-Series S4810 S4820T ... 453
vlt-peer-lag port-channel ... 454
Z-Series S4810 S4820T ... 454
Overview ... 454
VLT on Core Switches ... 455
Enhanced VLT ... 456
VLT Terminology ... 456
Configure Virtual Link Trunking ... 457
Important Points to Remember ... 457
Configuration Notes ... 458
Primary and Secondary VLT Peers ... 461
VLT Bandwidth Monitoring ... 462
VLT and Stacking ... 462
VLT and IGMP Snooping ... 462
VLT IPv6 ... 462
VLT Port Delayed Restoration ... 463
PIM-Sparse Mode Support on VLT ... 463
VLT Routing ... 465
Non-VLT ARP Sync ... 467
Verifying a VLT Configuration ... 467
Additional VLT Sample Configurations ... 471
Configuring Virtual Link Trunking (VLT Peer 1)Configuring Virtual Link Trunking (VLT Peer 2)Verifying a Port-Channel Connection to a VLT Domain (From an Attached Access Switch) ... 471
Troubleshooting VLT ... 473
FC Flex IO Modules ... 475
36 Understanding and Working of the FC Flex IO Modules ... 477

FC Flex IO Modules Overview..... 477
FC Flex IO Module Capabilities and Operations.....478
Guidelines for Working with FC Flex IO Modules.....479
Port Numbering for FC Flex IO Modules..... 480
Installing the Optics..... 481
Processing of Data Traffic.....481
Operation of the FIP Application..... 481
Operation of the NPIV Proxy Gateway..... 482
Installing and Configuring the Switch.....482
Installing and Configuring Flowchart for FC Flex IO Modules.....483
Installation.....484
Unpacking the Switch.....484
Interconnectivity of FC Flex IO Modules with Cisco MDS Switches..... 485

37 Data Center Bridging (DCB) for FC Flex IO Modules.....487

Interworking of DCB Map With DCB Buffer Threshold Settings.....487
dcb-map..... 488
priority-pgid..... 489
priority-group bandwidth pfc..... 490
dcb-map stack-unit all stack-ports all..... 491
show qos dcb-map.....492
DCB Command..... 493
dcb-enable.....493
DCBX Commands.....493
advertise dcbx-appln-tlv..... 494
advertise dcbx-tlv..... 494
dcbx port-role.....495
dcbx version.....496
debug dcbx..... 496
fcoe priority-bits..... 497
iscsi priority-bits.....498
show interface dcbx detail..... 498
ETS Commands.....501
bandwidth-percentage..... 501
clear ets counters..... 502
dcb-map.....502
dcb-output.....503
dcb-policy output.....504
dcb-policy output stack-unit stack-ports all..... 504
description.....505
ets mode on..... 505
priority-group..... 506

priority-group bandwidth pfc....507
priority-group qos-policy....508
priority-list....509
qos-policy-output ets....509
scheduler....510
set-pgid....511
show interface ets....511
show qos dcb-output....515
show qos priority-groups....515
show stack-unit stack-ports ets details....516
PFC Commands....517
clear pfc counters....517
dcb stack-unit pfc-buffering pfc-port-count pfc-queues....517
dcb-input....518
dcb-policy input....519
dcb-policy input stack-unit stack-ports all....520
description....520
pfc link-delay....521
pfc mode on....521
pfc no-drop queues....522
pfc priority....523
show dcb....523
show interface pfc....524
show interface pfc statistics....527
show qos dcb-input....527
show stack-unit stack-ports pfc details....528

38 Data Center Bridging (DCB)....529
Ethernet Enhancements in Data Center Bridging....529
Priority-Based Flow Control....530
Enhanced Transmission Selection....531
Configuring DCB Maps and its Attributes....533
Data Center Bridging: Default Configuration....536
Configuring PFC and ETS in a DCB Map....536
Applying a DCB Map in a Switch Stack....539
Data Center Bridging Exchange Protocol (DCBx)....539
Data Center Bridging in a Traffic Flow....540
Enabling Data Center Bridging....540
QoS dot1p Traffic Classification and Queue Assignment....541
Configure Enhanced Transmission Selection....542
ETS Operation with DCBx....542
Configuring Bandwidth Allocation for DCBx CIN....543

Configure a DCBx Operation....544
DCBx Operation....544
DCBx Port Roles....544
DCB Configuration Exchange....546
Configuration Source Election....546
Propagation of DCB Information....547
Auto-Detection and Manual Configuration of the DCBx Version....547
DCBx Example....548
DCBx Prerequisites and Restrictions....548
Configuring DCBx....549
Verifying the DCB Configuration....553
PFC and ETS Configuration Examples....564
Using PFC and ETS to Manage Data Center Traffic....564
PFC and ETS Configuration Command Examples....566
Using PFC and ETS to Manage Converged Ethernet Traffic in a Switch Stack....567
Hierarchical Scheduling in ETS Output Policies....567

39 Fibre Channel over Ethernet for FC Flex IO Modules....569

40 NPIV Proxy Gateway for FC Flex IO Modules....571
dcb-map....571
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....571
description (for FCoE maps)....572
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....572
fabric....572
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....572
fabric-id vlan....573
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....573
fcf-priority....574
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....574
fc-map....575
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....575
fcoe priority-bits....576
fcoe-map....576
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....576
fka-adv-period....577
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....577
interface vlan (NPIV proxy gateway)....578
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....578
keepalive....579
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....579
priority-group bandwidth pfc....579

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....579
show fcoe-map....581
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....581
show npiv devices....583
M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module....583
NPIV Proxy Gateway Configuration on FC Flex IO Modules....586
NPIV Proxy Gateway Operations and Capabilities....586
NPIV Proxy Gateway Operation....587
NPIV Proxy Gateway: Protocol Services....587
NPIV Proxy Gateway Functionality....588
NPIV Proxy Gateway: Terms and Definitions....588
Configuring an NPIV Proxy Gateway....590
Enabling Fibre Channel Capability on the Switch....590
Creating a DCB map....590
Applying a DCB map on server-facing Ethernet ports....592
Creating an FCoE VLAN....592
Creating an FCoE map....592
Applying an FCoE map on server-facing Ethernet ports....593
Applying an FCoE Map on fabric-facing FC ports....594
Sample Configuration....595
Displaying NPIV Proxy Gateway Information....595
show interfaces status Command Example....596
show fcoe-map Command Examples....597
show qos dcb-map Command Examples....598
show npiv devices brief Command Example....598
show npiv devices Command Example....599
show fc switch Command Example....600

1

About this Document

This document describes the new functionalities and enhancements in the Dell Networking OS Release version 9.3.0.0. All of the behavioral-changes and new features are covered in this single, consolidated Addendum. Use this document in conjunction with the hardware and software manuals of Release 9.2.0.0, which contains comprehensive information on the working and usage of the different platforms and their associated functionalities. You can obtain a copy of the latest documents of Release 9.2.0.0 from the technical documentation website at http://www.dell.com/manuals

We are not publishing the entire documentation set for Release version 9.3.0.0. Instead, this document presents the new and changed hardware and software processes for this release. It supplements the Release version 9.2.0.0 set of documents and allows you to locate information in an easy, streamlined way.

For topics that highlight the syntax and usage of commands, only the parameters that have been introduced or modified from the previous release are included in this document. However, the newly introduced commands, are however, covered in depth. For a complete description of all commands that have been enhanced or modified in Release 9.3.0.0 and were present in Release 9.2.0.0, refer the respective *Command Line Reference Guide* of the applicable platform.

For topics that provide a conceptual overview of new functionalities, and configuration procedures, only the enhancements and changes that have been implemented in Release 9.3.0.0 are mentioned in this Addendum. For complete information about such features that have been only enhanced and are not newly introduced in this release, refer the respective *Configuration Guide* of the applicable platform of Release 9.2.0.0.

NOTE: Although information that describes functionalites on the S4810 and S4820T platforms is included in this document, Dell Networking OS Release 9.3(0.0) is not supported on the S4810 and S4820T platforms.

Audience

This document is intended for system administrators who are responsible for configuring and maintaining networks and assumes knowledge in Layer 2 and Layer 3 networking technologies.

Conventions

This guide uses the following conventions to describe command syntax.

Keyword	Keywords are in Courier (a monospaced font) and must be entered in the CLI as listed.
parameter	Parameters are in italics and require a number or word to be entered in the CLI.
{X}	Keywords and parameters within braces must be entered in the CLI.
[X]	Keywords and parameters within brackets are optional.
x|y	Keywords and parameters separated by a bar require you to choose one option.

x||y	Keywords and parameters separated by a double bar allows you to choose any or all of the options.

Related Documents

For more information about the Dell Networking S4810, S4820T, S6000, Z9000, MXL 10/40GbE Switch, and the I/O Aggregator systems, refer the following documents corresponding to each of the platforms:

- *FTOS Command Reference*
- *Installing the System*
- *Dell Quick Start Guide*
- *FTOS Release Notes*

2

802.1X on the MXL 10/40GbE Switch

In Dell Networking OS Release 9.3(0.0), the MXL 10/40GbE Switch supports 802.1X port authentication. 802.1X is a method of port security. A device connected to a port that is enabled with 802.1X is disallowed from sending or receiving packets on the network until its identity can be verified (through a username and password, for example). For details on the command syntaxes and the keywords, refer the 802.1X chapter of the *MXL Command Reference Guide* of Release 9.2(0.2). For details on the conceptual overview and step-wise procedures to enable and configure 802.1X settings, refer the 802.1X chapter of the *MXL Configuration Guide* of Release 9.2(0.2).

3

ACL VLAN Groups and Content Addressable Memory (CAM)

This chapter describes the ACL VLAN group and CAM enhancements, and contains the following sections:

- Optimizing CAM Utilization During the Attachment of ACLs to VLANs
- Allocating FP Blocks for VLAN Processes

Optimizing CAM Utilization During the Attachment of ACLs to VLANs

This functionality is supported on the S4810, S4820T, Z9000, I/O Aggregator and MXL platforms.

You can enable and configure the access control list (ACL) content addressable memory (CAM) optimization functionality to minimize the number of entries in CAM while ACLs are applied on a VLAN or a set of VLANs and also while ACLS are applied on a set of ports. This capability enables effective usage of CAM space when Layer 3 ACLs are applied to a set of VLANs and when Layer 2 or Layer 3 ACLs are applied on a set of ports.

In releases of Dell Networking OS that does not support the CAM optimization functionality to reduce the usage of CAM area for application of ACLs, when an ACL is applied on a VLAN, the rules of the ACL are configured in the ACL region with the rule-specific parameters along with the Vlan as additional attributes. Therefore, when the ACL is applied on multiple VLAN interfaces, the consumption of CAM area increases proportionally. For example, when an ACL with 'n' number of rules is applied on 'm' number of VLAN interfaces, totally (n*m) entries are configured in the CAM region that is allocated for ACLs. Similarly, when an L2 or L3 ACL is applied on a set of ports, the same problem with large usage of CAM area occurs because a port is used as a parameter to be saved in CAM.

To avoid this problem of excessive consumption of CAM area, you can configure ACL VLAN groups that combines all the VLANs that are applied with the same ACL in a single group. A class identifier (Class ID) for each of ACL attached to the VLAN is assigned and this Class ID is used as an identifier or locator in the CAM area instead of the VLAN id. This method of processing signficiantly reduces the number of entries in the CAM area and saves memory space by using the class ID as filtering criterion in CAM instead of the VLAN ID.

You can create an ACL VLAN group and attach the ACL with the VLAN members. The optimization is applicable only when you create an ACL VLAN group. If you apply an ACL separately on the VLAN interface, each ACL has a mapping with the VLAN and increased CAM space utilization occurs. Attaching an ACL individually to VLAN interfaces is similar to the behavior of ACL-VLAN mapping storage in CAM prior to the implementation of the ACL VLAN group functionality.

The ACL manager application on router processor (RP1) contains all the state information about all the Acl Vlan groups that are present. The ACL handler on control processor (CP) and the ACL agent on line cards do not contain any stateful information about the group. The ACL manager application performs all the validation after you enter an `acl-vlan-group` command. If the command is valid, it is processed and sent to the agent if required. If a configuration error is found or if the maximum limit is exceeded for

the ACL VLAN groups present on the system, an appropriate error message is displayed. The ACL manager application processes the following parameters when you enter an `acl-vlan-group` command:

- Whether the CAM profile is set in VFP
- Whether the maximum number of groups in the system is exceeded
- Whether the maximum number of VLAN numbers permitted per ACL group is exceeded
- When a VLAN member that is being added is already a part of another ACL group

After these verification steps are performed, the ACL manager considers the command as valid and sends the information to the ACL agent on the line card as applicable. The ACL manager notifies the ACL agent in the following cases:

- A VLAN member is added or removed from a group and previously associated VLANs exist in the group
- Egress ACL is applied or removed from the group and the group contains VLAN members VLAN members are added or deleted from a vlan, which itself is a group member.
- A line card returns to the active state after going down and this line card contains a VLAN that is a member of an ACL group
- The ACL VLAN group is deleted and it contains VLAN members

The ACL manager does not notify the ACL agent in the following cases:

- The ACL VLAN group is created.
- The ACL VLAN group is deleted and it does not contain any VLAN members.
- The ACL is applied or removed from a group, and the ACL group does not contain a VLAN member.
- The description of the ACL group is added or removed.

Guidelines for Configuring ACL VLAN groups

This functionality is supported on the S4810, S4820T, Z9000, I/O Aggregator and MXL platforms.

Keep the following points in mind when you configure ACL VLAN groups:

- The interfaces to which the ACL VLAN group are applied function as restricted interfaces. The ACL VLAN group name is used to identify the group of VLANs that is used to perform hierarchical filtering.
- You can add only one ACL to an interface at a point in time.
- When you attempt to attach an ACL VLAN group to the same interface, a validation is performed to determine whether an ACL is applied directly to an interface. If you previously applied an ACL separately to the interface, an error occurs when you attempt to attach an ACL VLAN group to the same interface.
- The limitation on the maximum number of members that can be part of the ACL VLAN group is determined by the type of switch and its hardware capabilities. This scaling limit depends on the number of slices that are allocated for ACL CAM optimization. If one slice is allocated, the maximum number of VLAN members is 256 for all ACL VLAN groups. If two slices are allocated, the maximum number of VLAN members is 512 for all ACL VLAN groups.
- The maximum number of VLAN groups that you can configure also depends on the hardware specifications of the switch. Each VLAN group is mapped to a unique ID in the hardware. The maximum number of ACL VLAN groups supported is 31. Only a maximum two components (iSCSI

counters, Open Flow, ACL optimization) can be allocated virtual flow processing slices at a point in time.

- The maximum number of VLANs that you can configure as a member of ACL VLAN groups is limited to 512 on the S4180, Z9000, and MXL switches if two slices are allocated. If only one virtual flow processing slice is allocated, the maximum number of VLANs that you can configure as a member of an ACL VLAN group is 256 for the S4810, Z9000, and MXL switches.
- Port ACL optimization is applicable only for ACLs that are applied without the VLAN range.
- You cannot view the statistical details of ACL rules per VLAN and per interface if you enable the ACL VLAN group capability because this type of statistical information is available only for ACLs that are separately applied to VLANs. You can view the counters per ACL only.
- To display information using a particular ACL name, although you cannot display this detail using a specified interface name, you can use the `show ip accounting access list` command.
- Within a port, you can apply Layer 2 ACLs on a VLAN or a set of VLANs. In this case, CAM optimization is not applied.
- To enable optimization of CAM space for Layer 2 or Layer 3 ACLs that are applied to ports, the port number is removed as a qualifier for ACL application on ports, and port bits are used. When you apply the same ACL to a set of ports, the port bitmap is set when the ACL flow processor entry is added. When you remove the ACL from a port, the port bitmap is removed.
- If you do not attach an ACL to any of the ports, the flow processor entries are deleted. In this manner, when the same ACL is applied on set of ports, only one set of entries is installed in the flow processor (FP), thereby effectively saving CAM space. The optimization is enabled only if you specify the optimized option with the `ip access-group` command. This option is not valid for VLAN and LAG interfaces.

Configuring ACL VLAN Groups and Configuring FP Blocks for VLAN Parameters

. This section contains the following topics that describe how to configure ACL VLAN groups that you can attach to VLAN interfaces to optimize the utilization of CAM blocks and also how to configure flow processor (FP) blocks for different VLAN operations.

Configuring ACL VLAN Groups

You can create an ACL VLAN group and attach the ACL with the VLAN members. The optimization is applicable only when you create an ACL VLAN group. If you apply an ACL separately on the VLAN interface, each ACL has a mapping with the VLAN and increased CAM space utilization occurs. Attaching an ACL individually to VLAN interfaces is similar to the behavior of ACL-VLAN mapping storage in CAM prior to the implementation of the ACL VLAN group functionality.

1. Create an ACL VLAN group
   CONFIGURATION mode
   `acl-vlan-group {group name}`
   You can have up to eight different ACL VLAN groups at any given time.
2. Add a description to the ACL VLAN group.
   CONFIGURATION (conf-acl-vl-grp) mode
   `description description`

3. Apply an egress IP ACL to the ACL VLAN group.
   CONFIGURATION (conf-acl-vl-grp) mode
   ```
   ip access-group {group name} out implicit-permit
   ```
4. Add VLAN member(s) to an ACL VLAN group.
   CONFIGURATION (conf-acl-vl-grp) mode
   ```
   member vlan {VLAN-range}
   ```
5. Display all the ACL VLAN Groups or display a specific ACL VLAN Group, identified by name.
   CONFIGURATION (conf-acl-vl-grp) mode
   ```
   show acl-vlan-group {group name | detail}
   ```
   ```
   Dell#show acl-vlan-group detail

   Group Name :
     TestGroupSeventeenTwenty
   Egress IP Acl :
     SpecialAccessOnlyExpertsAllowed
   Vlan Members :
     100,200,300

   Group Name :
     CustomerNumberIdentificationEleven
   Egress IP Acl :
     AnyEmployeeCustomerElevenGrantedAccess
   Vlan Members :
     2-10,99

   Group Name :
     HostGroup
   Egress IP Acl :
     Group5
   Vlan Members :
     1,1000
   Dell#
   ```

Configuring FP Blocks for VLAN Parameters

You can use the `cam-acl-vlan` command to allocate the number of FP blocks for the various VLAN processes on the system. You can use the **no** version of this command to reset the number of FP blocks to default. By default, 0 groups are allocated for the ACL in VCAP. ACL VLAN groups or CAM optimization is not enabled by default, and you need to allocate the slices for CAM optimization.

1. Allocate the number of FP blocks for VLAN Open Flow operations.
   CONFIGURATION mode
   ```
   cam-acl-vlan vlanopenflow <0-2>
   ```
2. Allocate the number of FP blocks for VLAN iSCSI counters.
   CONFIGURATION mode
   ```
   cam-acl-vlan vlaniscsi <0-2>
   ```
3. Allocate the number of FP blocks for ACL VLAN optimization feature.
   CONFIGURATION mode
   ```
   cam-acl-vlan vlanaclopt <0-2>
   ```

4. View the number of flow processor (FP) blocks that is allocated for the different VLAN services.

   EXEC Privilege mode

```
Dell#show cam-usage switch

Linecard|Portpipe| CAM Partition   | Total CAM   |  Used CAM   |Available
CAM
========|========|=================|=============|=============|
==============
   11   |   0    | IN-L2 ACL       |      7152   |          0  |
7152
        |        | IN-L2 FIB       |     32768   |      1081   |
31687
        |        | OUT-L2 ACL      |         0   |          0  |
0
   11   |   1    | IN-L2 ACL       |      7152   |          0  |
7152
        |        | IN-L2 FIB       |     32768   |      1081   |
31687
        |        | OUT-L2 ACL      |         0   |          0  |         0
```

Viewing CAM Usage

This functionality is supported on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms.

View the amount of CAM space available, used, and remaining in each partition (including IPv4Flow and Layer 2 ACL sub- partitions) using the `show cam-usage` command from EXEC Privilege mode

Display Layer 2, Layer 3, ACL, or all CAM usage statistics.

EXCE Privilege mode

```
show cam usage [acl | router | switch]
```

The following sample output shows the consumption of CAM blocks for Layer 2 and Layer 3 ACLs, in addition to other processes that use CAM space:

```
Dell#show cam-usage
Linecard|Portpipe| CAM Partition   | Total CAM   |  Used CAM   |Available CAM
========|========|=================|=============|=============|==============
    1   |   0    | IN-L2 ACL       |      1008   |       320   |       688
        |        | IN-L2 FIB       |     32768   |      1132   |     31636
        |        | IN-L3 ACL       |     12288   |         2   |     12286
        |        | IN-L3 FIB       |    262141   |        14   |    262127
        |        | IN-L3-SysFlow   |      2878   |        45   |      2833
        |        | IN-L3-TrcList   |      1024   |         0   |      1024
        |        | IN-L3-McastFib  |      9215   |         0   |      9215
        |        | IN-L3-Qos       |      8192   |         0   |      8192
        |        | IN-L3-PBR       |      1024   |         0   |      1024
        |        | IN-V6 ACL       |         0   |         0   |         0
        |        | IN-V6 FIB       |         0   |         0   |         0
        |        | IN-V6-SysFlow   |         0   |         0   |         0
        |        | IN-V6-McastFib  |         0   |         0   |         0
        |        | OUT-L2 ACL      |      1024   |         0   |      1024
        |        | OUT-L3 ACL      |      1024   |         0   |      1024
        |        | OUT-V6 ACL      |         0   |         0   |         0
    1   |   1    | IN-L2 ACL       |       320   |         0   |       320
        |        | IN-L2 FIB       |     32768   |      1136   |     31632
        |        | IN-L3 ACL       |     12288   |         2   |     12286
        |        | IN-L3 FIB       |    262141   |        14   |    262127
        |        | IN-L3-SysFlow   |      2878   |        44   |      2834
--More--
```

The following sample output displays the CAM space utilization when Layer 2 and Layer 3 ACLs are configured:

```
Dell#show cam-usage acl
Linecard|Portpipe| CAM Partition    | Total CAM   |  Used CAM   |Available CAM
========|========|==================|=============|=============|============
   11   |   0    | IN-L2 ACL        |      1008   |          0  |      1008
        |        | IN-L3 ACL        |     12288   |          2  |     12286
        |        | OUT-L2 ACL       |      1024   |          2  |      1022
        |        | OUT-L3 ACL       |      1024   |          0  |      1024
```

The following sample output displays the CAM space utilization for Layer 2 ACLs:

```
Dell#show cam-usage switch

Linecard|Portpipe| CAM Partition    | Total CAM   |  Used CAM   |Available CAM
========|========|==================|=============|=============|==============
   11   |   0    | IN-L2 ACL        |      7152   |          0  |      7152
        |        | IN-L2 FIB        |     32768   |       1081  |     31687
        |        | OUT-L2 ACL       |         0   |          0  |         0
   11   |   1    | IN-L2 ACL        |      7152   |          0  |      7152
        |        | IN-L2 FIB        |     32768   |       1081  |     31687
        |        | OUT-L2 ACL       |         0   |          0  |         0
```

The following sample output displays the CAM space utilization for Layer 3 ACLs:

```
Dell#show cam-usage router
Linecard|Portpipe| CAM Partition    | Total CAM   |  Used CAM   |Available CAM
========|========|==================|=============|=============|==============
   11   |   0    | IN-L3 ACL        |      8192   |          3  |      8189
        |        | IN-L3 FIB        |    196607   |          1  |    196606
        |        | IN-L3-SysFlow    |      2878   |          0  |      2878
        |        | IN-L3-TrcList    |      1024   |          0  |      1024
        |        | IN-L3-McastFib   |      9215   |          0  |      9215
        |        | IN-L3-Qos        |      8192   |          0  |      8192
        |        | IN-L3-PBR        |      1024   |          0  |      1024
        |        | OUT-L3 ACL       |     16384   |          0  |     16384
   11   |   1    | IN-L3 ACL        |      8192   |          3  |      8189
        |        | IN-L3 FIB        |    196607   |          1  |    196606
        |        | IN-L3-SysFlow    |      2878   |          0  |      2878
        |        | IN-L3-TrcList    |      1024   |          0  |      1024
        |        | IN-L3-McastFib   |      9215   |          0  |      9215
        |        | IN-L3-Qos        |      8192   |          0  |      8192
        |        | IN-L3-PBR        |      1024   |          0  |      1024
        |        | OUT-L3 ACL       |     16384   |          0  |     16384
```

Allocating FP Blocks for VLAN Processes

This functionality is supported on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms.

The VLAN ContentAware Processor (VCAP) application is a pre-ingress CAP that modifies the VLAN settings before packets are forwarded. To support the ACL CAM optimization functionality, the CAM carving feature is enhanced. A total of four VACP groups are present, of which two are for fixed groups and the other two are for dynamic groups. Out of the total of two dynamic groups, you can allocate zero, one, or two FP blocks to iSCSI Counters, OpenFlow and ACL Optimization.

You can configure only two of these features at a point in time.

- To allocate the number of FP blocks for VLAN open flow operations, use the `cam-acl-vlan vlanopenflow <0-2>` command.

- To allocate the number of FP blocks for VLAN iSCSI counters , use the `cam-acl-vlan vlaniscsi <0-2>` command.
- To allocate the number of FP blocks for ACL VLAN optimization feature, use the `cam-acl-vlan vlanaclopt <0-2>` command.

You can use the **no** version of these commands to reset the number of FP blocks to default. By default, 0 groups are allocated for the ACL in VCAP. ACL VLAN groups or CAM optimization is not enabled by default, and you need to allocate the slices for CAM optimization.

To display the number of FP blocks that is allocated for the different VLAN services, you can use the `show cam-acl-vlan` command. After CAM configuration for ACL VLAN groups is performed, you must reboot the system to enable the settings to be stored in nonvolatile storage. During the initialization of CAM, the chassis manager reads the NVRAM and allocates the dynamic VCAP regions.

member vlan

Add VLAN members to an ACL VLAN group.

Syntax

```
member vlan {VLAN-range}
```

Parameters

VLAN-range	Enter the member VLANs using comma-separated VLAN IDs, a range of VLAN IDs, a single VLAN ID, or a combination. For example: Comma-separated: 3, 4, 6 Range: 5-10 Combination: 3, 4, 5-10, 8

Default None

Command Modes CONFIGURATION (conf-acl-vl-grp)

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms

Usage Information

At a maximum, there can be only 32 VLAN members in all ACL VLAN groups. A VLAN can belong to only one group at any given time.

You can create an ACL VLAN group and attach the ACL with the VLAN members. The optimization is applicable only when you create an ACL VLAN group. If you apply an ACL separately on the VLAN interface, each ACL has a mapping with the VLAN and increased CAM space utilization occurs.

Attaching an ACL individually to VLAN interfaces is similar to the behavior of ACL-VLAN mapping storage in CAM prior to the implementation of the ACL VLAN group functionality.

ip access-group

Apply an egress IP ACL to the ACL VLAN group.

Syntax	`ip access-group {group name} out implicit-permit`	
Parameters	***group-name***	Enter the name of the ACL VLAN group where you want the egress IP ACLs applied, up to 140 characters.
	out	Enter the keyword `out` to apply the ACL to outgoing traffic.
	implicit-permit	Enter the keyword `implicit-permit` to change the default action of the ACL from implicit-deny to implicit-permit (that is, if the traffic does not match the filters in the ACL, the traffic is permitted instead of dropped).
Default	None	
Command Modes	CONFIGURATION (conf-acl-vl-grp)	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms
Usage Information	You can apply only an egress IP ACL on an ACL VLAN group.	

show acl-vlan-group

Display all the ACL VLAN groups or display a specific ACL VLAN group, identified by name.

Syntax	`show acl-vlan-group {group-name | detail}`	
Parameters	***group-name***	(Optional) Display only the ACL VLAN group that is specified, up to 140 characters.
	detail	Display information in a line-by-line format to display the names in their entirety. Without the detail option, the output displays in a table style and information may be truncated.
Default	No default behavior or values	
Command Modes	EXEC EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms

Usage Information

When an ACL-VLAN-Group name or the Access List Group Name contains more than 30 characters, the name is truncated in the show `acl-vlan-group` command output.

Examples

The following sample illustrates the output of the `show acl-vlan-group` command.

NOTE: Some group names and some access list names are truncated.

```
Dell#show acl-vlan-group
Group Name                  Egress IP Acl               Vlan
Members
TestGroupSeventeenTwenty    SpecialAccessOnlyExperts
100,200,300
CustomerNumberIdentifica    AnyEmployeeCustomerEleve    2-10,99
HostGroup                   Group5                      1,1000

Dell#
```

The following sample output is displayed when using the `show acl-vlan-group` *`group-name`* option.

NOTE: The access list name is truncated.

```
Dell#show acl-vlan-group TestGroupSeventeenTwenty
Group Name                  Egress IP Acl               Vlan
Members
TestGroupSeventeenTwenty    SpecialAccessOnlyExperts
100,200,300

Dell#
```

The following sample output shows the line-by-line style display when using the `show acl-vlan-group detail` option.

NOTE: No group or access list names are truncated

```
Dell#show acl-vlan-group detail

Group Name :
  TestGroupSeventeenTwenty
Egress IP Acl :
  SpecialAccessOnlyExpertsAllowed
Vlan Members :
  100,200,300

Group Name :
  CustomerNumberIdentificationEleven
Egress IP Acl :
  AnyEmployeeCustomerElevenGrantedAccess
Vlan Members :
  2-10,99

Group Name :
  HostGroup
Egress IP Acl :
  Group5
Vlan Members :
  1,1000
```

```
Dell#
```

show cam-acl-vlan

Display the number of flow processor (FP) blocks that is allocated for the different VLAN services.

Syntax `show cam-acl-vlan`

Command Modes EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, Z9000, I/O Aggregator and MXL platforms.

Usage Information After CAM configuration for ACL VLAN groups is performed, you must reboot the system to enable the settings to be stored in nonvolatile storage. During the initialization of CAM, the chassis manager reads the NVRAM and allocates the dynamic VCAP regions.

The following table describes the output fields of this `show` command:

Field	Description
Chassis Vlan Cam ACL	Details about the CAM blocks allocated for ACLs for various VLAN operations at a system-wide, global level.
Stack Unit <number>	Details about the CAM blocks allocated for ACLs for various VLAN operations for a particular stack unit.
Current Settings(in block sizes)	Information about the number of FP blocks that are currently in use or allocated.
VlanOpenFlow	Number of FP blocks for VLAN open flow operations.
VlanIscsi	Number of FP blocks for VLAN internet small computer system interface (iSCSI) counters.
VlanHp	Number of FP blocks for VLAN high performance processes.
VlanFcoe	Number of FP blocks for VLAN Fiber Channel over Ethernet (FCoE) operations.
VlanAclOpt	Number of FP blocks for ACL VLAN optimzation feature.

Example

```
Dell#show cam-acl-vlan
-- Chassis Vlan Cam ACL --
             Current Settings(in block sizes)
```

```
VlanOpenFlow :          0
VlanIscsi    :          2
VlanHp       :          1
VlanFcoe     :          1
VlanAclOpt :          0

-- Stack unit 0 --
           Current Settings(in block sizes)
VlanOpenFlow :          0
VlanIscsi    :          2
VlanHp       :          1
VlanFcoe     :          1
VlanAclOpt :          0
```

cam-acl-vlan

Allocate the number of flow processor (FP) blocks or entries for VLAN services and processes.

Syntax

```
cam-acl-vlan { default | vlanopenflow <0-2> | vlaniscsi <0-2> |
vlanaclopt <0-2>
```

Parameters

default	Reset the number of FP blocks to default. By default, 0 groups are allocated for the ACL in VCAP. ACL VLAN groups or CAM optimization is not enabled by default, and you need to allocate the slices for CAM optimization.
vlanopenflow <0-2>	Allocate the number of FP blocks for VLAN open flow operations.
vlaniscsi <0-2>	Allocate the number of FP blocks for VLAN iSCSI counters.
vlanaclopt <0-2>	Allocate the number of FP blocks for the ACL VLAN optimization feature.

Default

If you use the `default` keyword with the `cam-acl-vlan` command, the FP blocks allocated for VLAN processes are restored to their default values. No FP blocks or dynamic VLAN ContentAware Processor (VCAP) groups are allocated for VLAN operations by default.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the S4810 and Z9000 platforms.

Usage Information

The VLAN ContentAware Processor (VCAP) application is a pre-ingress CAP that modifies the VLAN settings before packets are forwarded. To support the ACL CAM optimization functionality, the CAM carving feature is enhanced. A total of four VACP groups are present, of which two are for fixed groups and the other two are for dynamic groups. Out of the total of two dynamic groups, you can allocate zero, one, or two flow processor (FP) blocks to iSCSI Counters, OpenFlow and ACL Optimization. You can configure only two of these features at a point in time.

show cam-usage

View the amount of CAM space available, used, and remaining in each partition (including IPv4Flow and Layer 2 ACL sub-partitions).

Syntax

```
show cam-usage [acl | router | switch]
```

Parameters

acl	(OPTIONAL) Enter the keyword `acl` to display Layer 2 and Layer 3 ACL CAM usage.
router	(OPTIONAL) Enter the keyword `router` to display Layer 3 CAM usage.
switch	(OPTIONAL) Enter the keyword `switch` to display Layer 2 CAM usage.

Command Modes

EXEC

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, Z9000, I/O Aggregator and MXL platforms.

Usage Information

The following regions must be provided in the `show cam-usage` output:

- L3AclCam
- L2AclCam
- V6AclCam

The following table describes the output fields of this `show` command:

Field	Description
LineCard	Number of the line card that contains information on ACL VLAN groups
Portpipe	The hardware path that packets follow through a system for ACL optimization
CAM-Region	Type of area in the CAM block that is used for ACL VLAN groups
Total CAM space	Total amount of space in the CAM block
Used CAM	Amount of CAM space that is currently in use
Available CAM	Amount of CAM space that is free and remaining to be allocated for ACLs

Example 1: Output of the show cam-usage Command

```
Dell#show cam-usage
Linecard|Portpipe| CAM Partition   | Total CAM   |  Used CAM
|Available CAM
========|========|=================|=============|
=============|==============
    1   |   0    | IN-L2 ACL       |      1008   |       320
|       688
        |        | IN-L2 FIB       |     32768   |      1132
|     31636
        |        | IN-L3 ACL       |     12288   |         2
|     12286
        |        | IN-L3 FIB       |    262141   |        14
|    262127
        |        | IN-L3-SysFlow   |      2878   |        45
|      2833
        |        | IN-L3-TrcList   |      1024   |         0
|      1024
        |        | IN-L3-McastFib  |      9215   |         0
|      9215
        |        | IN-L3-Qos       |      8192   |         0
|      8192
        |        | IN-L3-PBR       |      1024   |         0
|      1024
        |        | IN-V6 ACL       |         0   |         0
|         0
        |        | IN-V6 FIB       |         0   |         0
|         0
        |        | IN-V6-SysFlow   |         0   |         0
|         0
        |        | IN-V6-McastFib  |         0   |         0
|         0
        |        | OUT-L2 ACL      |      1024   |         0
|      1024
        |        | OUT-L3 ACL      |      1024   |         0
|      1024
        |        | OUT-V6 ACL      |         0   |         0
|         0
    1   |   1    | IN-L2 ACL       |       320   |         0
|       320
        |        | IN-L2 FIB       |     32768   |      1136
|     31632
        |        | IN-L3 ACL       |     12288   |         2
|     12286
        |        | IN-L3 FIB       |    262141   |        14
|    262127
        |        | IN-L3-SysFlow   |      2878   |        44
|      2834
--More--
```

Example 2: Output of the show cam-usage acl Command

```
Dell#show cam-usage acl
Linecard|Portpipe| CAM Partition   | Total CAM   |  Used CAM
|Available CAM
========|========|=================|=============|
=============|============
   11   |   0    | IN-L2 ACL       |      1008   |
0    |      1008
        |        | IN-L3 ACL       |     12288   |
2    |     12286
        |        | OUT-L2 ACL      |      1024   |
2    |      1022
```

```
        |        | OUT-L3 ACL      |      1024   |
0    |      1024
```

Example 3: Output of the show cam-usage router Command

```
Dell#show cam-usage router
Linecard|Portpipe| CAM Partition   | Total CAM   |  Used CAM
|Available CAM
========|========|=================|=============|
=============|==============
   11   |   0    | IN-L3 ACL       |      8192   |         3
|      8189
        |        | IN-L3 FIB       |    196607   |         1
|    196606
        |        | IN-L3-SysFlow   |      2878   |         0
|      2878
        |        | IN-L3-TrcList   |      1024   |         0
|      1024
        |        | IN-L3-McastFib  |      9215   |         0
|      9215
        |        | IN-L3-Qos       |      8192   |         0
|      8192
        |        | IN-L3-PBR       |      1024   |         0
|      1024
        |        | OUT-L3 ACL      |     16384   |         0
|     16384
   11   |   1    | IN-L3 ACL       |      8192   |         3
|      8189
        |        | IN-L3 FIB       |    196607   |         1
|    196606
        |        | IN-L3-SysFlow   |      2878   |         0
|      2878
        |        | IN-L3-TrcList   |      1024   |         0
|      1024
        |        | IN-L3-McastFib  |      9215   |         0
|      9215
        |        | IN-L3-Qos       |      8192   |         0
|      8192
        |        | IN-L3-PBR       |      1024   |         0
|      1024
        |        | OUT-L3 ACL      |     16384   |         0
|     16384
```

Example 4: Output of the show cam-usage switch Command

```
Dell#show cam-usage switch

Linecard|Portpipe| CAM Partition   | Total CAM   |  Used CAM
|Available CAM
========|========|=================|=============|
=============|==============
   11   |   0    | IN-L2 ACL       |      7152   |         0
|      7152
        |        | IN-L2 FIB       |     32768   |      1081
|     31687
        |        | OUT-L2 ACL      |         0   |         0
|          0
   11   |   1    | IN-L2 ACL       |      7152   |         0
|      7152
        |        | IN-L2 FIB       |     32768   |      1081
|     31687
        |        | OUT-L2 ACL      |         0   |         0
|          0
```

show running config acl-vlan-group

Display the running configuration of all or a given ACL VLAN group.

Syntax `show running config acl-vlan-group` *group name*

Parameters

group-name	Display only the ACL VLAN group that is specified. The maximum group name is 140 characters.

Default None

Command Modes EXEC
EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms

Examples The following sample output shows the line-by-line style display when using the `show running-config acl-vlan-group` option. Note that no group or access list names are truncated

```
Dell#show running-config acl-vlan-group
!
acl-vlan-group group1
 description Acl Vlan Group1
 member vlan 1-10,400-410,500
 ip access-group acl1 out implicit-permit
!
acl-vlan-group group2
 member vlan 20
 ip access-group acl2 out
Dell#

Dell#show running-config acl-vlan-group group1
!
acl-vlan-group group1
 description Acl Vlan Group1
 member vlan 1-10,400-410,500
 ip access-group acl1 out implicit-permit

Dell#
```

acl-vlan-group

Create an ACL VLAN group.

Syntax `acl-vlan-group {`*group name*`}`

To remove an ACL VLAN group, use the `no acl-vlan-group {`*group name*`}` command.

Parameters

group-name	Specify the name of the ACL VLAN group. The name can contain a maximum 140 characters.

Default	No default behavior or values	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms

Usage Information

You can have up to eight different ACL VLAN groups at any given time. When you configure an ACL VLAN group, you enter the ACL VLAN Group Configuration mode.

To avoid the problem of excessive consumption of CAM area, you can configure ACL VLAN groups that combines all the VLANs that are applied with the same ACL in a single group. A unique identifier for each of ACL attached to the VLAN is used as a handle or locator in the CAM area instead of the VLAN id. This method of processing signficiantly reduces the number of entries in the CAM area and saves memory space in CAM.

You can create an ACL VLAN group and attach the ACL with the VLAN members. Optimization is applicable only when you create an ACL VLAN group. If you apply an ACL separately on the VLAN interface, each ACL maps with the VLAN and increased CAM space utilization occurs.

Attaching an ACL individually to VLAN interfaces is similar to the behavior of ACL-VLAN mapping storage in CAM prior to the implementation of the ACL VLAN group functionality.

show acl-vlan-group detail

Display all the ACL VLAN Groups or display a specific ACL VLAN Group by name. To display the names in their entirety, the output displays in a line-by-line format.

Syntax

```
show acl-vlan-group detail
```

Parameters

detail	Display information in a line-by-line format to display the names in their entirety. Without the detail option, the output is displayed in a table style and information may be truncated.

Default	No default behavior or values	
Command Modes	EXEC EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms

Usage Information

The output for this command displays in a line-by-line format. This allows the ACL-VLAN-Group names (or the Access List Group Names) to display in their entirety.

Examples The following sample output shows the line-by-line style display when using the `show acl-vlan-group detail` option. Note that no group or access list names are truncated

```
Dell#show acl-vlan-group detail

Group Name :
  TestGroupSeventeenTwenty
Egress IP Acl :
  SpecialAccessOnlyExpertsAllowed
Vlan Members :
  100,200,300

Group Name :
  CustomerNumberIdentificationEleven
Egress IP Acl :
  AnyEmployeeCustomerElevenGrantedAccess
Vlan Members :
  2-10,99

Group Name :
  HostGroup
Egress IP Acl :
  Group5
Vlan Members :
  1,1000
Dell#
```

description (ACL VLAN Group)

Add a description to the ACL VLAN group.

Syntax `description description`

Parameters

description	Enter a description to identify the ACL VLAN group (80 characters maximum).

Default No default behavior or values

Command Modes CONFIGURATION (conf-acl-vl-grp)

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms

Usage Information Enter a description for each ACL VLAN group that you create for effective and streamlined administrative and logging purposes.

4

Access Control Lists

This chapter describes the access control list (ACL) enhancements and contains the following sections:

- Logging of ACL Processes

Logging of ACL Processes

This functionality is supported on the S4810, S4820T, Z9000, I/O Aggregator and MXL platforms.

To assist in streamlined, robust administration and management of traffic that traverses the device after being validated by the configured ACLs, you can enable the generation of logs for access control list (ACL) processes. Although you can configure ACLs with the required permit or deny filters to provide access to the incoming packet or disallow access to a particular user, it is also necessary to monitor and examine the traffic that passes through the device. To enable such a mechanism to evaluate network traffic that is subjected to ACLs, you can configure the logs to be triggered for ACL operations. This functionality is primarily needed for network supervision and maintenance activities of the handled subscriber traffic.

If you configure logging of ACL activities, when a frame reaches an interface that is applied with an ACL and a match occurs against that ACL, that is installed with logging enabled, then whenever a frame that arrives at an interface hits a specific ACL entry, a log is generated to indicate details about the ACL entry that matched the packet.

A packet floe through a network path is defined by the source and destination IP addresses, protocols, and ports. Because the source port might differ for a new link between the same two hosts, instead of the same flow being used, a new flow might be created.

When you enable the generation of ACL log messages, at times, depending on the volume of traffic, it is possible that a large number of logs might be generated that can impact the system performance and efficiency. To avoid a storm of ACL logs from being recorded, you can configure a rate-limiting functionality to safeguard the system from an avalanche of ACL logs. You can specify the interval or frequency at which ACL logs must be triggered and also the threshold or the limit for the maximum number of logs to be generated. If you do not specify the frequency at which ACL logs must be generated, a default interval of 5 minutes is used. Similarly, if you do not specify the threshold for ACL logs, a default threshold of 10 is used, where this value refers to the number of packets that are matched against an ACL .

A Layer 2 or Layer 3 ACL contains a set of defined rules that are saved as flow processor (FP) entries. When you enable ACL logging for a particular ACL rule, a set of specific ACL rules translate to a set of FP entries. You can enable logging for each of these FP entries separately, which relates to each of the ACL entries configured in an ACL. For each ACL entry, the Dell Networking OS saves a table that maps each ACL entry that matches the received packet with the ACL name, sequence number of the rule, and the interface index in the database. When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

The ACL application sends the ACL logging configuration information and other details, such as the action, sequence number, and the ACL parameters that pertains to that ACL entry. The ACL service collects the ACL log records and records the following attributes per log message.

- For non-IP packets, the ACL name, sequence number, ACL action (permit or deny), source and destination MAC addresses, EtherType, and ingress interface are the logged attributes.
- For IP Packets, the ACL name, sequence number, ACL action (permit or deny), source and destination MAC addresses, source and destination IP addresses, the transport layer protocol used are the logged attributes.
- For IP packets that contain the transport layer protocol as Transmission Control Protocol (TCP) or User Datagram Protocol (UDP), the ACL name, sequence number, ACL action (permit or deny), source and destination MAC addresses, source and destination IP addresses, and the source and destination port (which are Layer 4 parameters) are also recorded.

If the packet contains an unidentified EtherType or transport layer protocol, the values for these parameters is saved as Unknown in the log message. If you also enable the count of packets for the ACL entry for which you configured logging, and if the logging is deactivated in a specific interval owing to the threshold being exceeded, the count of packets that exceeded the logging threshold value during that interval is logged when the subsequent log record is generated for that ACL entry in a different window or interval.

Guidelines for Configuring ACL Logging

This functionality is supported on the S4810, S4820T, Z9000, I/O Aggregator and MXL platforms.

Keep the following points in mind when you configure logging of ACL activities:

- During initialization, the ACL logging application tags the ACL rule indices for which a match condition exists as being in-use, which ensures that the same rule indices are not reused by ACL logging again.
- The ACL configuration information that the ACL logging application receives from the ACL manager causes the allocation and clearance of the match rule number. A unique match rule number is created for the combination of each ACL entry, sequence number, and interface parameters.
- A separate set of match indices is preserved by the ACL logging application for the permit and deny actions. Depending on the action of an ACL entry, the corresponding match index is allocated from the particular set that is maintained for permit and dent actions.
- The maximum number of ACL entries with permit action that can be logged is 125. The maximum number of ACL entries with deny action that can be logged is 126.
- For virtual ACL entries, the same match rule number is reused. Similarly, when an ACL entry is deleted that was previously enabled for ACL logging, the match rule number used by it is released back to the pool or set of match indices that is present so that it can be reused for subsequent allocations.
- The ACL logging application saves the allocated match rule number in the ACL entry itself so that it can be reused when the ACL entry is reprogrammed due to CAM changes.
- The allocated match rule number for an ACL entry is associated with an FP entry and saved in the system. A timer control starts when an FP entry is added to the system or CPU with the logging option, and the timer stops when the ACL entry is deleted. The ACL logger module obtains the ACL name, sequence number, and interface index from the match rule index contained in the packet.
- A maximum of 15 ACL entries or records can be saved in the space that is allocated for ACL logging.

- A timer control of 30 seconds is present in the ACL agent module, the expiry of which causes the log records that are collocted until that time are transmitted to the ACL manager for logging. An inter-process communication (IPC) message is sent to the ACL manager by the ACL agent when a maximum of 15 records are collected or the 30-second timer period is exceeded.
- If you enabled the count of packets for the ACL entry for which you configured logging, and if the logging is deactivated in a specific interval owing to the threshold being exceeded, the count of packets that exceeded the logging threshold value during that interval is logged when the subsequent log record is generated for that ACL entry in a different window or interval.
- When you delete an ACL entry, the logging settings associated with it are also removed.
- ACL logging is supported for standard and extended IPv4 ACLs, IPv6 ACLs, and standard and extended MAC ACLs.
- For ACL entries applied on port-channel interfaces, one match index for every member interface of the port-channel interface is assigned. Therefore, the total available match indices of 251 are split (125 match indices for permit action and 126 match indices for the deny action).
- You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.
- The total uniquely available match rule indices is 255 with four match indices used by other modules, leaving 51 indices available for ACL logging.

Configuring ACL Logging

This functionality is supported on the S4810, S4820T, Z9000, I/O Aggregator and MXL platforms.

To configure the maximum number of ACL log messages to be generated and the frequency at which these messages must be generated, perform the following:

NOTE: This example describes the configuration of ACL logging for standard IP access lists. You can enable the logging capability for standard and extended IPv4, IPv6, and standard and extended MAC ACLs.

1. Specify the maximum number of ACL logs or the threshold that can be generated by using the `threshold-in-msgs` *`count`* option with the seq, permit, or deny commands. Upon exceeding the specified maximum limit, the generation of ACL logs is terminated. You can enter a threshold in the range of 1-100. By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

   CONFIG-STD-NACL mode

   ```
   seq sequence-number {deny | permit} {source [mask] | any | host ip-address}
   [count [byte]] [order] [fragments] [log [threshold-in-msgs count] ]
   ```

2. Specify the interval in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes. The default frequency at which ACL logs are generated is 5 minutes. If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, IPv6 ACLs, and standard and extended MAC ACLs. Configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

   CONFIG-STD-NACL mode

   ```
   seq sequence-number {deny | permit} {source [mask] | any | host ip-address}
   [count [byte]] [order] [fragments] [log [interval minutes]]
   ```

deny (for Standard IP ACLs)

To drop packets with a certain IP address, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`deny {source | any | host {ip-address}}[count [byte]] [dscp value] [order] [fragments] [log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny {source [mask] | any | host ip-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-STANDARD-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard

and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

ip access-list standard — configures a standard ACL.

permit — configures a permit filter.

deny (for Extended IP ACLs)

Configure a filter that drops IP packets meeting the filter criteria.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
deny {ip | ip-protocol-number} {source mask | any | host ip-address} {destination mask | any | host ip-address} [count [byte]] [dscp value] [order] [monitor] [fragments] [log [interval minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny {ip | ip-protocol-number} {source mask | any | host ip-address} {destination mask | any | host ip-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny tcp — assigns a filter to deny TCP packets.

deny udp — assigns a filter to deny UDP packets.

ip access-list extended — creates an extended ACL.

seq (for Standard IPv4 ACLs)

Assign a sequence number to a deny or permit filter in an IP access list while creating the filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`seq` *`sequence-number`* `{deny | permit}` `{`*`source`* `[`*`mask`*`] | any | host` *`ip-address`*`}} [count [bytes]] [dscp` *`value`*`] [order] [fragments] [log [interval` *`minutes`*`] [threshold-in-msgs [count]]`

To delete a filter, use the `no seq` *`sequence-number`* command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-STANDARD-ACCESS-LIST

Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny — configures a filter to drop packets.

permit — configures a filter to forward packets.

deny tcp (for Extended IP ACLs)

Configure a filter that drops transmission control protocol (TCP) packets meeting the filter criteria.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
deny tcp {source mask | any | host ip-address} [bit] [operator
port [port]] {destination mask | any | host ip-address} [dscp]
[bit] [operator port [port]] [count [byte]] [order] [fragments]
[log [interval minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny tcp {source mask | any | host ip-address} {destination mask | any | host ip-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.

	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny — assigns a filter to deny IP traffic.

deny udp — assigns a filter to deny UDP traffic.

deny udp (for Extended IP ACLs)

To drop user datagram protocol (UDP) packets meeting the filter criteria, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
deny udp {source mask | any | host ip-address} [operator port
[port]] {destination mask | any | host ip-address} [dscp]
[operator port [port]] [count [byte]] [order] [fragments] [log
[interval minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny udp {source mask | any | host ip-address} {destination mask | any | host ip-address}` command.

Parameters	**log**	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
	threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.
Defaults	By default, 10 ACL logs are generated if you do not specify the threshold explicitly. The default frequency at which ACL logs are generated is 5 minutes.	
Command Modes	CONFIGURATION-EXTENDED-ACCESS-LIST	
Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.
Usage Information	When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval. If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.	
Related Commands	deny — assigns a filter to deny IP traffic. deny tcp — assigns a filter to deny TCP traffic.	

deny arp (for Extended MAC ACLs)

Configure an egress filter that drops ARP packets on egress ACL supported line cards. (For more information, refer to your line card documentation).

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
deny arp {destination-mac-address mac-address-mask | any} vlan
vlan-id {ip-address | any | opcode code-number} [count [byte]]
[order] [log [interval minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny arp {destination-mac-address mac-address-mask | any} vlan vlan-id {ip-address | any | opcode code-number}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

deny icmp (for Extended IP ACLs)

To drop all or specific internet control message protocol (ICMP) messages, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`deny icmp {source mask | any | host ip-address} {destination mask | any | host ip-address} [dscp] [message-type] [count [byte]] [order] [fragments] [log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny icmp {source mask | any | host ip-address} {destination mask | any | host ip-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

deny ether-type (for Extended MAC ACLs)

Configure an egress filter that drops specified types of Ethernet packets on egress ACL supported line cards. (For more information, refer to your line card documentation).

Syntax

`deny ether-type` *`protocol-type-number`* {*`destination-mac-address mac-address-mask`* | `any`} `vlan` *`vlan-id`* {*`source-mac-address mac-address-mask`* | `any`} `[count [byte]] [order] [log [interval` *`minutes`*`] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq` *`sequence-number`* command if you know the filter's sequence number.
- Use the `no deny ether-type` *`protocol-type-number`* {*`destination-mac-address mac-address-mask`* | `any`} `vlan` *`vlan-id`* {*`source-mac-address mac-address-mask`* | `any`} command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are

applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

deny (for Standard MAC ACLs)

To drop packets with a the MAC address specified, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`deny {any | mac-source-address [mac-source-address-mask]} [count [byte]] [log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny {any | mac-source-address mac-source-address-mask}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-MAC ACCESS LIST-STANDARD

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for

standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

permit — configures a MAC address filter to pass packets.

seq — configures a MAC address filter with a specified sequence number.

deny (for Extended MAC ACLs)

To drop packets that match the filter criteria, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`deny {any | host mac-address | mac-source-address mac-source-address-mask} {any | host mac-address | mac-destination-address mac-destination-address-mask} [ethertype-operator] [count [byte]][log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny {any | host mac-address | mac-source-address mac-source-address-mask} {any | host mac-address | mac-destination-address mac-destination-address-mask}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-MAC ACCESS LIST-EXTENDED

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

permit — configures a MAC address filter to pass packets.

seq — configures a MAC address filter with a specified sequence number.

permit arp (for Extended MAC ACLs)

Configure a filter that forwards ARP packets meeting this criteria. This command is supported only on 12-port GE line cards with SFP optics; refer to your line card documentation for specifications.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
permit arp {destination-mac-address mac-address-mask | any}
vlan vlan-id {ip-address | any | opcode code-number} [count
[byte]] [order] [fragments] [log [interval minutes] [threshold-
in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `{destination-mac-address mac-address-mask | any} vlan vlan-id {ip-address | any | opcode code-number}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes	CONFIGURATION-EXTENDED-ACCESS-LIST	
Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

permit ether-type (for Extended MAC ACLs)

Configure a filter that allows traffic with specified types of Ethernet packets. This command is supported only on 12-port GE line cards with SFP optics. For specifications, refer to your line card documentation.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
permit ether-type protocol-type-number {destination-mac-address
mac-address-mask | any} vlan vlan-id {source-mac-address mac-
address-mask | any} [count [byte]] [order] [log [interval
minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no permit ether-type protocol-type-number {destination-mac-address mac-address-mask | any} vlan vlan-id {source-mac-address mac-address-mask | any}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.

	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

permit icmp (for Extended IP ACLs)

Configure a filter to allow all or specific ICMP messages.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`permit icmp {`*`source mask`*` | any | host `*`ip-address`*`} {`*`destination mask`*` | any | host `*`ip-address`*`} [dscp] [`*`message-type`*`] [count [byte]] [order] [fragments] [log [interval `*`minutes`*`] [threshold-in-msgs [`*`count`*`]]`

To remove this filter, you have two choices:

- Use the `no seq `*`sequence-number`* command if you know the filter's sequence number.
- Use the `no permit icmp {`*`source mask`*` | any | host `*`ip-address`*`} {`*`destination mask`*` | any | host `*`ip-address`*`}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.

	threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-STANDARD-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

permit udp (for Extended IP ACLs)

To pass UDP packets meeting the filter criteria, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`permit udp {`*`source mask`*` | any | host `*`ip-address`*`} [`*`operator port`* `[`*`port`*`]] {`*`destination mask`*` | any | host `*`ip-address`*`} [dscp] [`*`operator port`* `[`*`port`*`]] [count [byte]] [order] [fragments] [log [interval `*`minutes`*`] [threshold-in-msgs [`*`count`*`]]`

To remove this filter, you have two choices:

- Use the `no seq `*`sequence-number`* command if you know the filter's sequence number.
- Use the `no permit udp {`*`source mask`*` | any | host `*`ip-address`*`} {`*`destination mask`*` | any | host `*`ip-address`* command.

Parameters	**log**	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
	threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.
Defaults	By default, 10 ACL logs are generated if you do not specify the threshold explicitly. The default frequency at which ACL logs are generated is 5 minutes.	
Command Modes	CONFIGURATION-EXTENDED-ACCESS-LIST	
Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.
Usage Information	When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval. If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.	
Related Commands	ip access-list extended — creates an extended ACL. permit — assigns a permit filter for IP packets. permit tcp — assigns a permit filter for TCP packets.	

permit (for Extended IP ACLs)

To pass IP packets meeting the filter criteria, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`permit {source mask | any | host ip-address} {destination mask | any | host ip-address} [count [bytes]] [dscp value] [order] [fragments] [log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no deny {source mask | any | host ip-address} {destination mask | any | host ip-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

ip access-list extended — creates an extended ACL.

permit tcp — assigns a permit filter for TCP packets.

permit udp — assigns a permit filter for UDP packets.

permit (for Standard MAC ACLs)

To forward packets from a specific source MAC address, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`permit {any | mac-source-address [mac-source-address-mask]} [count [byte]] | [log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no permit {any | mac-source-address mac-source-address-mask}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-MAC ACCESS LIST-STANDARD

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard

and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny — configures a MAC ACL filter to drop packets.

seq —configure a MAC ACL filter with a specified sequence number.

seq (for Standard MAC ACLs)

To a deny or permit filter in a MAC access list while creating the filter, assign a sequence number.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`seq sequence-number {deny | permit} {any | mac-source-address [mac-source-address-mask]} [count [byte]] [log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, use the `no seq sequence-number` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-MAC ACCESS LIST-STANDARD

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for

standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny — configures a filter to drop packets.

permit — configures a filter to forward packets.

permit tcp (for Extended IP ACLs)

To pass TCP packets meeting the filter criteria, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
permit tcp {source mask | any | host ip-address} [bit]
[operator port [port]] {destination mask | any | host ip-
address} [bit] [dscp] [operator port [port]] [count [byte]]
[order] [fragments] [log [interval minutes] [threshold-in-msgs
[count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no permit tcp {source mask | any | host ip-address} {destination mask | any | host ip-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

ip access-list extended — creates an extended ACL.

permit — assigns a permit filter for IP packets.

permit udp — assigns a permit filter for UDP packets.

seq arp (for Extended MAC ACLs)

Configure an egress filter with a sequence number that filters ARP packets meeting this criteria. This command is supported only on 12-port GE line cards with SFP optics. For specifications, refer to your line card documentation.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
seq sequence-number {deny | permit} arp {destination-mac-address mac-address-mask | any} vlan vlan-id {ip-address | any | opcode code-number} [count [byte]] [order] [log [interval minutes] [threshold-in-msgs [count]]
```

To remove this filter, use the no seq *sequence-number* command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes	CONFIGURATION-EXTENDED-ACCESS-LIST	
Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

seq ether-type (for Extended MAC ACLs)

Configure an egress filter with a specific sequence number that filters traffic with specified types of Ethernet packets. This command is supported only on 12-port GE line cards with SFP optics. For specifications, refer to your line card documentation.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
seq sequence-number {deny | permit} ether-type protocol-type-number {destination-mac-address mac-address-mask | any} vlan vlan-id {source-mac-address mac-address-mask | any} [count [byte]] [order] [log [interval minutes] [threshold-in-msgs [count]]
```

To remove this filter, use the no seq *sequence-number* command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

CONFIGURATION-EXTENDED-ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

seq (for IP ACLs)

Assign a sequence number to a deny or permit filter in an extended IP access list while creating the filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
seq sequence-number {deny | permit} {ip-protocol-number | icmp
| ip | tcp | udp} {source mask | any | host ip-address}
{destination mask | any | host ip-address} [operator port
[port]] [count [byte]] [dscp value] [order] [fragments] [log
[interval minutes] [threshold-in-msgs [count]]
```

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults	By default, 10 ACL logs are generated if you do not specify the threshold explicitly. The default frequency at which ACL logs are generated is 5 minutes.	
Command Modes	CONFIGURATION-EXTENDED-ACCESS-LIST	
Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.
Usage Information	When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval. If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.	
Related Commands	deny — configures a filter to drop packets. permit — configures a filter to forward packets.	

seq (for IPv6 ACLs)

Assign a sequence number to a deny or permit the filter in an IPv6 access list while creating the filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
seq sequence-number {deny | permit} {ipv6-protocol-number |
icmp | ip | tcp | udp} {source address mask | any | host ipv6-
address} {destination address | any | host ipv6-address}
[operator port [port]] [count [byte]] [log [interval minutes]
[threshold-in-msgs [count]]
```

To delete a filter, use the `no seq sequence-number` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.

	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny – configures a filter to drop packets.

permit – configures a filter to forward packets.

permit udp (for IPv6 ACLs)

Configure a filter to pass UDP packets meeting the filter criteria.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
permit udp {source address mask | any | host ipv6-address}
[operator port [port]] {destination address | any | host ipv6-
address} [operator port [port]] [count [byte]] [log [interval
minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no permit udp {source address mask | any | host ipv6-address} {destination address | any | host ipv6-address}` command.

Parameters	**log**	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
	threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.
Defaults	By default, 10 ACL logs are generated if you do not specify the threshold explicitly. The default frequency at which ACL logs are generated is 5 minutes.	
Command Modes	ACCESS-LIST	
Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.
Usage Information	When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval. If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.	
Related Commands	permit – assigns a permit filter for IP packets. permit tcp – assigns a permit filter for TCP packets.	

permit tcp (for IPv6 ACLs)

Configure a filter to pass TCP packets that match the filter criteria.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
permit tcp {source address mask | any | host ipv6-address}
[operator port [port]] {destination address | any | host ipv6-
```

address} [*bit*] [*operator port* [*port*]] [count [byte]] [log [interval *minutes*] [threshold-in-msgs [*count*]]

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no permit tcp {source address mask | any | host ipv6-address} {destination address | any | host ipv6-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

permit – assigns a permit filter for IP packets.

permit udp – assigns a permit filter for UDP packets.

permit icmp (for IPv6 ACLs)

To allow all or specific internet control message protocol (ICMP) messages, configure a filter.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`permit icmp {source address mask | any | host ipv6-address} {destination address | any | host ipv6-address} [message-type] [count [byte]] [log [interval minutes] [threshold-in-msgs [count]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command if you know the filter's sequence number.
- Use the `no permit icmp {source address mask | any | host ipv6-address} {destination address | any | host ipv6-address}` command.

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for

standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

permit (for IPv6 ACLs)

To configure a filter that matches the filter criteria, select an IPv6 protocol number, ICMP, IPv6, TCP, or UDP.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`permit {`*`ipv6-protocol-number`*` | icmp | ipv6 | tcp | udp} [count [byte]] [dscp `*`value`*`] [order] [fragments] [log [interval `*`minutes`*`] [threshold-in-msgs [`*`count`*`]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command syntax if you know the filter's sequence number
- Use the `no permit {`*`ipv6-protocol-number`*` | icmp | ipv6 | tcp | udp}` command

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that

new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

deny udp (for IPv6 ACLs)

Configure a filter to drop user datagram protocol (UDP) packets meeting the filter criteria.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

`deny udp {`*`source address mask`*` | any | host `*`ipv6-address`*`} [`*`operator port`*` [`*`port`*`]] {`*`destination address`*` | any | host `*`ipv6-address`*`} [`*`operator port`*` [`*`port`*`]] [count [byte]] [log [interval `*`minutes`*`] [threshold-in-msgs [`*`count`*`]]`

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command syntax if you know the filter's sequence number
- Use the `no deny udp {`*`source address mask`*` | any | host `*`ipv6-address`*`} {`*`destination address`*` | any | host `*`ipv6-address`*`}` command

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

ACCESS-LIST

Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny – assigns a filter to deny IP traffic.

deny tcp – assigns a deny filter for TCP traffic.

deny tcp (for IPv6 ACLs)

Configure a filter that drops TCP packets that match the filter criteria.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
deny tcp {source address mask | any | host ipv6-address}
[operator port [port]] {destination address | any | host ipv6-
address} [bit] [operator port [port]] [count [byte]] [log
[interval minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command syntax if you know the filter's sequence number
- Use the `no deny tcp {source address mask | any | host ipv6-address} {destination address | any | host ipv6-address}` command

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.

	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Related Commands

deny – assigns a filter to deny IP traffic.

deny udp – assigns a filter to deny UDP traffic.

deny icmp (for Extended IPv6 ACLs)

Configure a filter to drop all or specific ICMP messages.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
deny icmp {source address mask | any | host ipv6-address}
{destination address | any | host ipv6-address} [message-type]
[count [byte]] | [log]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command syntax if you know the filter's sequence number
- Use the `no deny icmp {source address mask | any | host ipv6-address} {destination address | any | host ipv6-address}` command

Parameters	**log**	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
	threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
	interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.
Defaults	By default, 10 ACL logs are generated if you do not specify the threshold explicitly. The default frequency at which ACL logs are generated is 5 minutes.	
Command Modes	ACCESS-LIST	
Command History	**Version 9.3.0.0**	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.
Usage Information	When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval. If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.	

deny (for IPv6 ACLs)

Configure a filter that drops IPv6 packets that match the filter criteria.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Syntax

```
deny {ipv6-protocol-number | icmp | ipv6 | tcp | udp} [count
[byte]] [dscp value] [order] [fragments] [log [interval
minutes] [threshold-in-msgs [count]]
```

To remove this filter, you have two choices:

- Use the `no seq sequence-number` command syntax if you know the filter's sequence number

- Use the `no deny {ipv6-protocol-number | icmp | ipv6 | tcp | udp}` command

Parameters

log	(OPTIONAL) Enter the keyword `log` to enable the triggering of ACL log messages.
threshold-in msgs *count*	(OPTIONAL) Enter the `threshold-in-msgs` keyword followed by a value to indicate the maximum number of ACL logs that can be generated, exceeding which the generation of ACL logs is terminated. with the seq, permit, or deny commands. You can enter a threshold in the range of 1-100.
interval *minutes*	(OPTIONAL) Enter the keyword `interval` followed by the time period in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes.

Defaults

By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

The default frequency at which ACL logs are generated is 5 minutes.

Command Modes

ACCESS-LIST

Command History

Version 9.3.0.0	Added support for logging of ACLs on the S4810, S4820T, Z9000, M I/O Aggregator and MXL 10/40GbE Switch IO Module platforms.

Usage Information

When the configured maximum threshold is exceeded, generation of logs is stopped. When the interval at which ACL logs are configured to be recorded expires, the subsequent, fresh interval timer is started and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold is exceeded, it is reenabled for this new interval.

If ACL logging is stopped because the configured threshold is exceeded, it is reenabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs. You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

Flow-Based Monitoring Support for ACLs

This functionality to enable flow-based monitoring is supported on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

Flow-based monitoring conserves bandwidth by monitoring only specified traffic instead all traffic on the interface. This feature is particularly useful when looking for malicious traffic. It is available for Layer 2 and Layer 3 ingresss traffic. You may specify traffic using standard or extended access-lists. This mechanism copies all incoming packets on one port and forwards (mirrors) them to another port. The source port is the monitored port (MD) and the destination port is the monitoring port (MG).

The port mirroring application maintains and performs all the monitoring operations on the chassis. ACL information is sent to the ACL manager, which in turn notifies the ACL agent to add entries in the CAM area. Duplicate entries in the ACL are not saved.

When a packet arrives at a port that is being monitored, the packet is validated against the configured ACL rules. If the packet matches an ACL rule, the system examines corresponding flow processor to perform the action specified for that port. If mirroring action is set in the flow processor entry, the destination port details, which indicates the port on the device to which the mirrored information must be sent, are sent to the destination port.

When a stack unit is reset or a stack unit undergoes a failure, the ACL agent registers with the port mirroring application. The port mirroring utility downloads the monitoring configuration to the ACL agent. The interface manager notifies the port mirroring application about the removal of an interface when an interface to which an ACL entry is associated is deleted.

Behavior of Flow-Based Monitoring

You can enter activate flow-based monitoring for a monitoring session by entering the `flow-based enable` command in the Monitor Session mode. When you enable this capability, traffic with particular flows that are traversing through the ingress interfaces are examined and, appropriate ACLs can be applied in the ingress direction. By default, flow-based monitoring is not enabled.

You must specify the monitor option with the `permit, deny`, or `seq` command for ACLs that are assigned to the source or the monitored port (MD) to enable the evaluation and replication of traffic that is destined to the source port to the destination port. Enter the keyword `monitor` with the `seq`, `permit` and `deny` ACL rules to allow or drop IPv4, IPv6, ARP, UDP, EtherType, ICMP, and TCP packets when the rule is describing the traffic that you want to monitor and the ACL in which you are creating the rule will be applied to the monitored interface. Flow monitoring is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs.

CONFIG-STD-NACL mode

`seq` *`sequence-number`* {`deny | permit`} {`source [mask] | any | host`} *`ip-address`*} `[count [byte]] [order] [fragments] [log [threshold-in-msgs` *`count`*`] ]`

If the number of monitoring sessions increases, inter-process communication (IPC) bandwidth utilization will be high. ACL manager might require a large bandwidth when you assign an ACL with many entries to an interface.

The ACL agent module saves monitoring details in its local database and also in the CAM region to monitor packets which match the specified criterion. The ACL agent maintains data on the source port, destination port, and the endpoint to which the packet must be forwarded when a match occurs with the ACL entry.

If you configure the `flow-based enable` command and do not apply an ACL on the source port or the monitored port, both flow-based monitoring and port mirroring do not function. Flow-based monitoring is supported only for ingress traffic and not for egress packets.

The port mirroring application maintains database that contains all monitoring sessions (including port monitor sessions). It has information regarding the sessions that are enabled for flow-based monitoring and those sessions that are not enabled for flow-based monitoring. It downloads monitoring configuration to the ACL agent whenever the ACL agent is registered with the port mirroring application or when flow-based monitoring is enabled.

The `show monitor session` *`session-id`* command has been enhanced to display the Type field in the output, which indicates whether a particular session is enabled for flow-monitoring.

Example Output of the `show` Command

```
E1200-maa-01#show running-config monitor session
!
monitor session 11
 flow-based enable
 source GigabitEthernet 13/0 destination GigabitEthernet 13/1 direction both
```

The `show running-config monitor session` displays whether flow-based monitoring is enabled for a particular session.

Example Output of the `show` Command

```
E1200-maa-01#show running-config monitor session
!
monitor session 11
 flow-based enable
 source GigabitEthernet 13/0 destination GigabitEthernet 13/1 direction both
```

The `show config` command has been modified to display monitoring configuration in particular session.

Example Output of the `show` Command

```
E1200-maa-01(conf-mon-sess-11)#show config
    !
    monitor session 11
                            flow-based enable
      source GigabitEthernet 13/0 destination GigabitEthernet 13/1 direction
both
```

The `show ip | mac | ipv6 accounting` commands have been enhanced to display whether monitoring is enabled for traffic that match with the rule of the specific ACL rules.

Example Output of the `show` Command

```
Force10# show ip accounting access-list
!
Extended Ingress IP access list kar on GigabitEthernet 10/0
Total cam count 1
   seq 5 permit ip 192.168.20.0/24 173.168.20.0/24 monitor

Force10#show mac accounting access-list kar in gi 10/0 out
Egress Extended mac access-list kar on GigabitEthernet 10/0
 seq 5  permit host  11:11:11:11:11:11 host  22:22:22:22:22:22 monitor
     seq 10  permit host  22:22:22:22:22:22 any monitor
     seq 15  permit host  00:0f:fe:1e:de:9b host  0a:0c:fb:1d:fc:aa monitor

Force10#show ipv6 accounting access-list
    !
    Ingress IPv6 access list kar on GigabitEthernet 10/0
    Total cam count 1
     seq 5 permit ipv6 22::/24 33::/24 monitor
```

Enabling Flow-Based Monitoring

Flow-based monitoring is supported on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

Flow-based monitoring conserves bandwidth by monitoring only specified traffic instead all traffic on the interface. This feature is particularly useful when looking for malicious traffic. It is available for Layer 2 and Layer 3 ingress and egress traffic. You may specify traffic using standard or extended access-lists.

1. Enable flow-based monitoring for a monitoring session.
   MONITOR SESSION mode
   `flow-based enable`
2. Define in access-list rules that include the keyword `monitor`. FTOS only considers for port monitoring traffic matching rules with the keyword `monitor`.
   CONFIGURATION mode
   `ip access-list`
   Refer to Access Control Lists (ACLs).
3. Apply the ACL to the monitored port.
   INTERFACE mode
   `ip access-group access-list`

To view an access-list that you applied to an interface, use the `show ip accounting access-list` command from EXEC Privilege mode.

Example of the `flow-based enable` Command

```
FTOS(conf)#monitor session 0
FTOS(conf-mon-sess-0)#flow-based enable
FTOS(conf)#ip access-list ext testflow
FTOS(config-ext-nacl)#seq 5 permit icmp any any count bytes monitor
FTOS(config-ext-nacl)#seq 10 permit ip 102.1.1.0/24 any count bytes monitor
FTOS(config-ext-nacl)#seq 15 deny udp any any count bytes
FTOS(config-ext-nacl)#seq 20 deny tcp any any count bytes
FTOS(config-ext-nacl)#exit
FTOS(conf)#interface gig 1/1
FTOS(conf-if-gi-1/1)#ip access-group testflow in
FTOS(conf-if-gi-1/1)#show config
!
interface GigabitEthernet 1/1
  ip address 10.11.1.254/24
  ip access-group testflow in
  shutdown
FTOS(conf-if-gi-1/1)#exit
FTOS(conf)#do show ip accounting access-list testflow
!
Extended Ingress IP access list testflow on GigabitEthernet 1/1
Total cam count 4
  seq 5 permit icmp any any monitor count bytes (0 packets 0 bytes)
  seq 10 permit ip 102.1.1.0/24 any monitor count bytes (0 packets 0 bytes)
  seq 15 deny udp any any count bytes (0 packets 0 bytes)
  seq 20 deny tcp any any count bytes (0 packets 0 bytes)
FTOS(conf)#do show monitor session 0
SessionID Source Destination Direction Mode      Type
--------- ------ ----------- --------- ----      ----
0         Gi 1/1 Gi 1/2      rx        interface Flow-based
```

5

Bare Metal Provisioning (BMP)

This chapter describes the Bare Metal Provisioning (BMP) enhancements that apply to the S4810, S4820T, S6000, Z9000, and MXL platforms

Support for BMP on the S6000 Switch

Starting with Dell Networking OS Release 9.3(0.0), BMP 3.1 is supported on the S6000 platform. For details about the commands and configuration procedures of BMP 3.1, refer the *Open Automation Guide*.

Enhanced Behavior of the stop bmp Command

The `stop bmp` command behaves as follows in different circumstances:

- While FTOS image upgrade is in-progress, aborts the BMP process once the FTOS image is upgraded.
- When applying configurations from file, aborts the BMP process after all configurations are applied in the system.
- When running pre-configuration or post-configuration scripts, stops execution of the script and aborts the BMP process immediately.
- While downloading the configuration or script file, aborts BMP process after download, neither applies configuration nor runs the script.

When you enter the CONFIGURATION mode during the BMP process, warning or error messages are displayed appropriately to avoid any configuration conflicts between user and the BMP process.

Removal of the Deprecated User-Defined String Parameter With reload-type Command

The `user-defined-string` parameter available with the `reload-type` command, which was deprecated in Dell Networking OS release 9.2(0.0) and earlier, is now removed. The `vendor-class-identifier` parameter replaces the `user-defined-string` parameter.

Inclusion of Service Tag Information in the Option 60 String

You can now configure the vendor class identifier up to a maximum of 128 characters. In the vendor class identifier (option 60) string, the User String field is also included with the Type, Hardware, Serial Number, Service Tag and OS Version fields.

Replacement of stop jump-start Command With the stop bmp Command

The deprecated `stop jump-start` command is replaced by the `stop bmp` from BMP 3.1 onwards. However, in BMP 1.5 and 2.0, you can use the `stop jump-start` command to stop the device from restarting in BMP mode.

6

Data Center Bridging (DCB)

This chapter describes the DCB enhancements and contains the following sections:

- Managing Hardware Buffer Statistics
- Configuring WRED and ECN Attributes
- Enabling Buffer Statistics Tracking
- Configuring DCB Maps and its Attributes
- Data Center Bridging: Default Configuration
- Configuring the Dynamic Buffer Method
- Priority-Based Flow Control Using Dynamic Buffer Method

Configuring DCB Maps and its Attributes

This topic contains the following sections that describe how to configure a DCB map, apply the configured DCB map to a port, configure PFC without a DCB map, and configure lossless queues. This functionality is supported S4810, S4820T, S6000, I/O Aggregator, and MXL platforms.

DCB Map: Configuration Procedure

A DCB map consists of PFC and ETS parameters. By default, PFC is not enabled on any 802.1p priority and ETS allocates equal bandwidth to each priority. To configure user-defined PFC and ETS settings, you must create a DCB map.

Step	Task	Command	Command Mode
1	Enter global configuration mode to create a DCB map or edit PFC and ETS settings.	dcb-map *name*	CONFIGURATION
2	Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group or whether priority group traffic should be handled with strict priority scheduling. You can enable PFC on a maximum of two priority queues on an interface. Enabling PFC for dot1p priorities makes the corresponding port queue lossless. The sum of all allocated bandwidth percentages in all groups in the DCB map must be 100%. Strict-priority traffic is serviced first. Afterwards, bandwidth allocated to other priority groups is made available and allocated according to the specified percentages. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups. Example: **priority-group 0 bandwidth 60 pfc off priority-group 1 bandwidth 20 pfc on**	priority-group *group_num* {**bandwidth** *percentage* | **strict-priority**} **pfc** {**on** | **off**}	DCB MAP

Step	Task	Command	Command Mode
	priority-group 2 bandwidth 20 pfc on priority-group 4 strict-priority pfc off Repeat this step to configure PFC and ETS traffic handling for each priority group.		
3	Specify the dot1p priority-to-priority group mapping for each priority. Priority-group range: 0 to 7. All priorities that map to the same queue must be in the same priority group. Leave a space between each priority group number. For example: **priority-pgid 0 0 0 1 2 4 4 4** in which priority group 0 maps to dot1p priorities 0, 1, and 2; priority group 1 maps to dot1p priority 3; priority group 2 maps to dot1p priority 4; priority group 4 maps to dot1p priorities 5, 6, and 7.	priority-pgid *dot1p0_group_num dot1p1_group_num dot1p2_group_num dot1p3_group_num dot1p4_group_num dot1p5_group_num dot1p6_group_num dot1p7_group_num*	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (**no priority pgid** command), the PFC and ETS parameters revert to their default values on the interfaces on which the DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.

  As a result, PFC and lossless port queues are disabled on 802.1p priorities, and all priorities are mapped to the same priority queue and equally share port bandwidth.
- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, first create a new DCB map with the desired PFC and ETS settings and apply the new map to the interfaces to override the previous DCB map settings. Then delete the original dot1p priority-priority group mapping.

  If you delete the dot1p priority-priority group mapping (**no priority pgid** command) before you apply the new DCB map, the default PFC and ETS parameters are applied on the interfaces. This change may create a DCB mismatch with peer DCB devices and interrupt network operation.

Applying a DCB Map on a Port

When you apply a DCB map with PFC enabled on an S6000 interface, a memory buffer for PFC-enabled priority traffic is automatically allocated. The buffer size is allocated according to the number of PFC-enabled priorities in the assigned map.

To apply a DCB map to an Ethernet port, follow these steps:

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	interface {**tengigabitEthernet** *slot/port* | **fortygigabitEthernet** *slot/port*}	CONFIGURATION
2	Apply the DCB map on the Ethernet port to configure it with the PFC and ETS settings in the map; for example: FTOS# **interface tengigabitEthernet 0/0**	dcb-map *name*	INTERFACE

Step	Task	Command	Command Mode
	FTOS(config-if-te-0/0)# **dcb-map SAN_A_dcb_map1** Repeat Steps 1 and 2 to apply a DCB map to more than one port. You cannot apply a DCB map on an interface which has been already configured for PFC using the **pfc priority** command or which is already configured for lossless queues (**pfc no-drop queues** command).		

Configuring PFC without a DCB Map

In a network topology that uses the default ETS bandwidth allocation (assigns equal bandwidth to each priority), you can also enable PFC for specified dot1p-priorities on individual interfaces without using a DCB map. This type of DCB configuration is useful on interfaces which require PFC for lossless traffic but do not transmit converged Ethernet traffic.

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	interface {tengigabitEthernet *slot*/***port*** | fortygigabitEthernet *slot*/***port***}	CONFIGURATION
2	Enable PFC on specified priorities. Range: 0-7. Default: None. Maximum number of loss less queues supported on an Ethernet port: 2. Separate priority values with a comma. Specify a priority range with a dash, for example: pfc priority 3,5-7 1. You cannot configure PFC using the **pfc priority** command on an interface on which a DCB map has been applied or which is already configured for lossless queues (**pfc no-drop queues** command).	pfc priority *priority-range*	INTERFACE

Configuring Lossless Queues

DCB also supports the manual configuration of lossless queues on an interface after you disable PFC mode in a DCB map and apply the map on the interface. The configuration of no-drop queues provides flexibility for ports on which PFC is not needed but lossless traffic should egress from the interface.

Lossless traffic egresses out the no-drop queues. Ingress 802.1p traffic from PFC-enabled peers is automatically mapped to the no-drop egress queues.

When configuring lossless queues on a port interface, take into account:

- By default, no lossless queues are configured on a port.
- A limit of two lossless queues are supported on a port. If the number of lossless queues configured exceeds the maximum supported limit per port (two), an error message displays. You must reconfigure the value to a smaller number of queues.
- If you configure lossless queues on an interface that already has a DCB map with PFC enabled (**pfc on**), an error message displays.

Step	Task	Command	Command Mode
1	Enter INTERFACE Configuration mode.	interface {**tengigabitEthernet** *slot/port* | **fortygigabitEthernet** *slot/port*}	CONFIGURATION
2	Open a DCB map and enter DCB map configuration mode.	dcb-map *name*	INTERFACE
3	Disable PFC.	no pfc mode on	DCB MAP
4	Return to interface configuration mode.	exit	DCB MAP
5	Apply the DCB map created to disable PFC operation on the interface	dcb-map {*name* | **default**}	INTERFACE
6	Configure the port queues that still function as no-drop queues for lossless traffic. For the dot1p-queue assignments, see Table 131. The maximum number of lossless queues globally supported on a port is 2. You cannot configure PFC no-drop queues on an interface on which a DCB map with PFC enabled has been applied or which is already configured for PFC using the **pfc priority** command. Range: 0-3. Separate queue values with a comma; specify a priority range with a dash; for example: pfc no-drop queues 1,3 or pfc no-drop queues 2-3 Default: No lossless queues are configured.	**pfc no-drop queues** *queue-range*	INTERFACE

Data Center Bridging: Default Configuration

This functionality is supported on the S6000 platform.

Before you configure PFC and ETS on an S5000 switch (see Configuring DCB Maps and its Attributes), take into account the following default settings:

DCB is enabled (see Enabling Data Center Bridging).

The PFC memory buffer supports up to 64 PFC-enabled ports and two lossless queues per port.

PFC and ETS are globally enabled by default:

The default dot1p priority-queue assignments are applied as follows:

802.1p value in incoming frame:	0	1	2	3	4	5	6	7
Egress queue assignment	0	0	0	1	2	3	3	3

PFC is not applied on specific dot1p priorities.

ETS: Equal bandwidth is assigned to each port queue and each dot1p priority in a priority group.

To configure PFC and ETS parameters on an S5000 interface, you must specify a PFC mode and ETS bandwidth allocation for a priority group and an 802.1p priority-to-priority group mapping in a DCB map (see Configuring PFC and ETS in a DCB Map). No default PFC and ETS settings are applied to Ethernet interfaces.

Configuring PFC and ETS in a DCB Map

An S6000 switch supports the use of a DCB map in which you configure priority-based flow control and enhanced transmission selection settings. To configure PFC and ETS parameters, you must apply a DCB map on an S6000 interface. This functionality is supported on the S6000 platform.

PFC Configuration Notes

Priority-based flow control (PFC) provides a flow control mechanism based on the 802.1p priorities in converged Ethernet traffic received on an interface and is enabled by default when you enable DCB. As an enhancement to the existing Ethernet pause mechanism, PFC stops traffic transmission for specified priorities (CoS values) without impacting other priority classes. Different traffic types are assigned to different priority classes.

When traffic congestion occurs, PFC sends a pause frame to a peer device with the CoS priority values of the traffic that needs to be stopped. DCBx provides the link-level exchange of PFC parameters between peer devices. PFC allows network administrators to create zero-loss links for SAN traffic that requires no-drop service, while at the same time retaining packet-drop congestion management for LAN traffic.

On an S6000 switch, PFC is enabled by default on Ethernet ports (**pfc mode on** command). You can configure PFC parameters using a DCB map or the **pfc priority** command in Interface configuration mode. For more information, see Configuring DCB Maps and its Attributes.

NOTE: DCB maps are supported only on physical Ethernet interfaces.

When you configure PFC in a DCB map:

- As soon as you apply a DCB map with PFC enabled on an interface, DCBx starts exchanging information with a peer. The IEEE802.1Qbb, CEE and CIN versions of PFC TLV are supported. DCBx also validates PFC configurations that are received in TLVs from peer devices.

By applying a DCB map with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, configure PFC priorities on all DCB egress ports.

- To remove a DCB map, including the PFC configuration it contains, use the **no dcb map** command in Interface configuration mode.
- To disable PFC operation on an interface, use the **no pfc mode on** command in DCB-Map configuration mode.
- Traffic may be interrupted when you reconfigure PFC no-drop priorities in a DCB map or re-apply the DCB map to an interface.
- For PFC to be applied, the configured priority traffic must be supported by a PFC peer (as detected by DCBx).
- If you apply a DCB map with PFC disabled (**pfc off**):
- You can enable link-level flow control on the interface (**flowcontrol rx on tx on** command; see Using Ethernet Pause Frames for Flow Control). To delete the DCB map, first disable link-level flow control. PFC is then automatically enabled on the interface because an interface is PFC-enabled by default.
- To ensure no-drop handling of lossless traffic, PFC allows you to configure lossless queues on a port (see Configuring DCB Maps and its Attributes).
- When you configure a DCB map, an error message displays if:
- The PFC dot1p priorities result in more than two lossless queues.

- When you apply a DCB map, an error message displays if:
- Link-level flow control is already enabled on an interface. You cannot enable PFC and link-level flow control at the same time on an interface.
- In a switch stack, configure all stacked ports with the same PFC configuration.
- FTOS allows you to change the default dot1p priority-queue assignments only if the change satisfies the following requirements in DCB maps already applied to S6000 interfaces:
- All 802.1p priorities mapped to the same queue must be in the same priority group.
- A maximum of two PFC-enabled, lossless queues are supported on an interface.

Otherwise the reconfiguration of a default dot1p-queue assignment is rejected.

- To ensure complete no-drop service, apply the same PFC parameters on all PFC-enabled peers.

PFC Prerequisites and Restrictions

On an S6000 switch, PFC is globally enabled by default, but not applied on specific 802.1p priorities. To enable PFC on 802.1p priorities, create a DCB map. For more information, see Configuring DCB Maps and its Attributes.

The following prerequisites and restrictions apply when you configure PFC in a DCB map:

- You can enable PFC on a maximum of two priority queues on an interface. Enabling PFC for dot1p priorities configures the corresponding port queue as lossless.
- You cannot enable PFC and link-level flow control at the same time on an interface.

ETS Configuration Notes

ETS provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic. Different traffic types have different service needs. Using ETS, you can create groups within an 802.1p priority class to configure different treatment for traffic with different bandwidth, latency, and best-effort needs.

When you configure ETS in a DCB map:

- The DCB map associates a priority group with a PFC operational mode (on or off) and an ETS scheduling and bandwidth allocation. You can apply a DCB map on multiple egress ports.
- Use the ETS configuration associated with 802.1p priority traffic in a DCB map in DCBx negotiation with ETS peers.
- Traffic in priority groups is assigned to strict-queue or weighted round-robin (WRR) scheduling in an ETS configuration and is managed using the ETS bandwidth-assignment algorithm. FTOS de-queues all frames of strict-priority traffic before servicing any other queues. A queue with strict-priority traffic can starve other queues in the same port.
- ETS-assigned bandwidth allocation and strict-priority scheduling apply only to data queues, not to control queues.
- FTOS supports hierarchical scheduling on an interface. FTOS control traffic is redirected to control queues as higher priority traffic with strict priority scheduling. After the control queues drain out, the remaining data traffic is scheduled to queues according to the bandwidth and scheduler configuration in the DCB map. The available bandwidth calculated by the ETS algorithm is equal to the link bandwidth after scheduling non-ETS higher-priority traffic.

- The configuration of bandwidth allocation and strict-queue scheduling is not supported at the same time for a priority group.
- **Bandwidth assignment**: By default, equal bandwidth is assigned to each dot1p priority in a priority group. To configure the bandwidth assigned to the port queues associated with dot1p priorities in a priority group, use the **bandwidth** *percentage* parameter. The sum of the bandwidth allocated to all priority groups in a DCB map must be 100% of the bandwidth on the link. You must allocate at least 1% of the total bandwidth to each priority group.
- **Scheduling of priority traffic**: dot1p priority traffic on the switch is scheduled to the current queue mapping. dot1p priorities within the same queue must have the same traffic properties and scheduling method.
- **ETS configuration error**: If an error occurs in an ETS configuration, the configuration is ignored and the scheduler and bandwidth allocation settings are reset to the ETS default value: 100% of available bandwidth is allocated to priority group 0 and bandwidth is equally assigned to each dot1p priority.

If an error occurs when a port receives a peer's ETS configuration, the port's configuration resets to the ETS configuration in the previously configured DCB map. If no DCB map was previously applied, the port resets to the default ETS parameters.

ETS Prerequisites and Restrictions

On an S6000 switch, ETS is enabled by default on Ethernet ports; equal bandwidth is assigned to each 802.1p priority. You can change the default ETS configuration only by using a DCB map. For more information, see Configuring DCB Maps and its Attributes.

The following prerequisites and restrictions apply when you configure ETS bandwidth allocation or strict-priority queuing in a DCB map:

- When allocating bandwidth or configuring strict-priority queuing for dot1p priorities in a priority group on a DCBx CIN interface, take into account the CIN bandwidth allocation (see Configuring Bandwidth Allocation for DCBx CIN) and dot1p-queue mapping.
- Although ETS bandwidth allocation or strict-priority queuing does not support weighted random early detection (WRED), explicit congestion notification (ECN), rate shaping, and rate limiting because these parameters are not negotiated by DCBx with peer devices, you can apply a QoS output policy with WRED and/or rate shaping on a DCBx CIN-enabled interface (see Configuring Port-based Rate Shaping and Weighted Random Early Detection). In this case, the WRED or rate shaping configuration in the QoS output policy must take into account the bandwidth allocation or queue scheduler configured in the DCB map.

Priority-Group Configuration Notes

When you configure priority groups in a DCB map:

- A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share the same latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.
- In a DCB map, each 802.1p priority must map to a priority group.
- The maximum number of priority groups supported in a DCB map on an interface is equal to the number of data queues (4) on the port. Each priority group can support more than one data queue.
- You can enable PFC on a maximum of two priority queues on an interface.

- If you configure more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

dcb-map

Create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic. Apply the DCB map to an Ethernet interface.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

`dcb-map` *`map-name`*

Parameters

map-name	Enter a DCB map name. The maximum number of alphanumeric characters is 32.

Defaults

None. There are no pre-configured PFC and ETS settings on S5000 Ethernet interfaces.

Command Modes

CONFIGURATION

INTERFACE

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

A DCB map is a template used to configure DCB parameters and apply them on converged Ethernet interfaces. DCB parameters include priority-based flow control (PFC) and enhanced traffic selection (ETS).

To display the PFC and ETS settings in DCB maps, enter the `show qos dcb-map` command.

Use the `dcb-map` command to create a DCB map to specify PFC and ETS settings and apply it on Ethernet ports. After you apply a DCB map to an interface, the PFC and ETS settings in the map are applied when the Ethernet port is enabled. DCBx is enabled on Ethernet ports by default.

The `dcb-map` command is supported only on physical Ethernet interfaces.

To remove a DCB map from an interface, enter the `no dcb-map` *`map-name`* command in Interface configuration mode.

Related Commands

show qos dcb-map– displays the dcb-map profiles configured on the system.

dcb-map stack-unit all stack-ports all– applies a DCB map on all ports of a switch stack.

priority-pgid

Assign 802.1p priority traffic to a priority group in a DCB map.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

```
priority-pgid dot1p0_group-num dot1p1_group-num dot1p2_group-num dot1p3_group-num dot1p4_group-num dot1p5_group-num dot1p6_group-num dot1p7_group-num
```

Parameters

dot1p0_group-num	Enter the priority group number for each 802.1p class of traffic in a DCB map.
dot1p1_group-num	
dot1p2_group-num	
dot1p3_group-num	
dot1p4_group-num	
dot1p5_group-num	
dot1p6_group-num	
dot1p7_group-num	

Defaults None

Command Modes DCB MAP

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

PFC and ETS settings are not pre-configured on Ethernet ports. You must use the `dcb-map` command to configure different groups of 802.1p priorities with PFC and ETS settings.

Using the `priority-pgid` command, you assign each 802.1p priority to one priority group. A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group. For example, the `priority-pgid 0 0 0 1 2 4 4 4` command creates the following groups of 802.1p priority traffic:

- Priority group 0 contains traffic with dot1p priorities 0, 1, and 2.
- Priority group 1 contains traffic with dot1p priority 3.
- Priority group 2 contains traffic with dot1p priority 4.
- Priority group 4 contains traffic with dot1p priority 5, 6, and 7.

To remove a priority-pgid configuration from a DCB map, enter the `no priority-pgid` command.

Related Commands

dcb-map — creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

priority-group bandwidth pfc— configures the ETS bandwidth allocation and the PFC setting used to manage the port traffic in an 802.1p priority group.

pfc mode on

Enable the PFC configuration on the port so that the priorities are included in DCBX negotiation with peer PFC devices.

Syntax

`pfc mode on`

To disable the PFC configuration, use the `no pfc mode on` command.

Defaults

PFC mode is on.

Command Modes

DCB INPUT POLICY

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

By applying a DCB input policy with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, also enable PFC on all DCB egress ports or configure the dot1p priority-queue assignment of PFC priorities to lossless queues (refer to `pfc no-drop queues`).

To disable PFC operation on an interface, enter the `no pfc mode on` command in DCB Input Policy Configuration mode. PFC is enabled and disabled as global DCB operation is enabled (`dcb-enable`) or disabled (`no dcb-enable`).

You cannot enable PFC and link-level flow control at the same time on an interface.

Related Commands

dcb-input — creates a DCB input policy.

priority-group bandwidth pfc

Configure the ETS bandwidth allocation and PFC mode used to manage port traffic in an 802.1p priority group.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

```
priority-group group-num {bandwidth percentage| strict-priority} pfc {on | off}
```

Parameters

priority-group *group-num*	Enter the keyword `priority-group` followed by the number of an 802.1p priority group. Use the `priority-pgid` command to create the priority groups in a DCB map.
bandwidth *percentage*	Enter the keyword `bandwidth` followed by a bandwidth percentage allocated to the priority group. The range of valid values is 1 to 100. The sum of all allocated bandwidth percentages in priority groups in a DCB map must be 100%.
strict-priority	Configure the priority-group traffic to be handled with strict priority scheduling. Strict-priority traffic is serviced first, before bandwidth allocated to other priority groups is made available.
pfc {on | off}	Configure whether priority-based flow control is enabled (on) or disabled (off) for port traffic in the priority group.

Defaults

None

Command Modes

DCB MAP

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

Use the `dcb-map` command to configure priority groups with PFC and/or ETS settings and apply them to Ethernet interfaces.

Use the `priority-pgid` command to map 802.1p priorities to a priority group. You can assign each 802.1p priority to only one priority group. A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.

Repeat the `priority-group bandwidth pfc` command to configure PFC and ETS traffic handling for each priority group in a DCB map.

You can enable PFC on a maximum of two priority queues.

If you configure more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.

To remove a priority-group configuration in a DCB map, enter the `no priority-group bandwidth pfc` command.

By default, equal bandwidth is assigned to each dot1p priority in a priority group. Use the `bandwidth` parameter to configure the bandwidth percentage assigned to a priority group. The sum of the bandwidth allocated to all priority groups in a DCB map must be 100% of the bandwidth on the link. You must allocate at least 1% of the total port bandwidth to each priority group.

Related Commands

dcb-map – creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

priority-pgid – configures the 802.1p priority traffic in a priority group for a DCB map.

dcb-map stack-unit all stack-ports all

Apply the specified DCB map on all ports of the switch stack.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

`dcb-map stack-unit all stack-ports all` *`dcb-map-name`*

To remove the PFC and ETS settings in a DCB map from all stack units, use the `no dcb-map stack-unit all stack-ports all` command.

Parameters

dcb-map-name	Enter the name of the DCB map.

Defaults

none

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

The `dcb-map stack-unit all stack-ports all` command overwrites any previous DCB maps applied to stack ports.

Related Commands

dcb-map – creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

show qos dcb-map

Display the DCB parameters configured in a specified DCB map.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

```
show qos dcb-map map-name
```

Parameters

map-name	Displays the PFC and ETS parameters configured in the specified map.

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

Use the `show qos dcb-map` command to display the enhanced transmission selection (ETS) and priority-based flow control (PFC) parameters used to configure server-facing Ethernet ports. S5000 Ethernet ports are DCBx-enabled by default.

The following table describes the `show qos dcb-map` output shown in the example below.

Field	Description
State	Complete: All mandatory DCB parameters are correctly configured. In progress: The DCB map configuration is not complete. Some mandatory parameters are not configured.
PFC Mode	PFC configuration in DCB map: On (enabled) or Off.
PG	Priority group configured in the DCB map.
TSA	Transmission scheduling algorithm used by the priority group: Enhanced Transmission Selection (ETS).
BW	Percentage of bandwidth allocated to the priority group.
PFC	PFC setting for the priority group: On (enabled) or Off.
Priorities	802.1p priorities configured in the priority group.

Example

```
FTOS# show qos dcb-map dcbmap2

State   :Complete
PfcMode:ON
--------------------
PG:0 TSA:ETS  BW:50  PFC:OFF
Priorities:0 1 2 4 5 6 7
```

```
PG:1 TSA:ETS  BW:50  PFC:ON
Priorities:3
```

Related Commands

dcb-map — creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

Priority-Based Flow Control Using Dynamic Buffer Method

Priority-based flow control using dynamic buffer spaces is supported on the S4810, S4820T, S6000, and MXL platforms.

In a data center network, priority-based flow control (PFC) manages large bursts of one traffic type in multiprotocol links so that it does not affect other traffic types and no frames are lost due to congestion. When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device.

Pause and Resume of Traffic

The pause message is a mechanism that is used by the sending device to inform the receiving device regarding a congested, heavily-loaded traffic state that has been identified. When the interface of a sending device transmits a pause frame, the recipient acknowledges this frame by temporarily halting the transmission of data packets. The sending device requests the recipient to restart the transmission of data traffic when the congestion eases and reduces. The time period that is specified in the pause frame defines the duration for which the flow of data packets is halted. When the time period elapses, the transmission restarts.

When a device sends a pause frame to another device, the time for which the sending of packets from the other device must be stopped is contained in the pause frame. The device that sent the pause frame empties the buffer to be less than the threshold value and restarts the acceptance of data packets.

Dynamic ingress buffering enables the sending of pause frames at different thresholds based on the number of ports that experience congestion at a point in time. This behavior impacts the total buffer size used by a particular lossless priority on an interface. The pause and resume thresholds can also be configured dynamically. You can configure a buffer size, pause threshold, ingress shared threshold weight, and resume threshold to control and manage the total amount of buffers that are to be used in your network environment.

All the PFC-related settings such as the DCB input and output policies or DCB maps are saved in the DCB application and the Differentiated Services Manager (DSM) application. All of these configurations can be modified only for interfaces that are enabled for DCB. The DCB buffer configurations are also saved in the DCB and DSM databases.

Buffer Sizes for Lossless or PFC Packets

You can configure up to a maximum of 4 lossless (PFC) queues. By configuring 4 lossless queues, you can configure 4 different priorities and assign a particular priority to each application that your network is used to process. For example, you can assign a higher priority for time-sensitive applications and a lower priority for other services, such as file transfers. You can configure the amount of buffer space to be allocated for each priority and the pause or resume thresholds for the buffer. This method of configuration enables you to effectively manage and administer the behavior of lossless queues.

Although the system contains 9 MB of space for shared buffers, a minimum guaranteed buffer is provided to all the internal and external ports in the system for both unicast and multicast traffic. This minimum

guaranteed buffer reduces the total available shared buffer to 7,787 KB. This shared buffer can be used for lossy and lossless traffic.

The default behavior causes up to a maximum of 6.6 MB to be used for PFC-related traffic. The remaining approximate space of 1 MB can be used by lossy traffic. You can allocate all the remaining 1 MB to lossless PFC queues. If you allocate in such a way, the performance of lossy traffic is reduced and degraded. Although you can allocate a maximum buffer size, it is used only if a PFC priority is configured and applied on the interface.

The number of lossless queues supported on the system is dependent on the availability of total buffers for PFC. The default configuration in the system guarantees a minimum of 52 KB per queue if all the 128 queues are congested. However, modifying the buffer allocation per queue impacts this default behavior.

By default the total available buffer for PFC is 6.6 MB and when you configure dynamic ingress buffering, a minimum of least 52 KB per queue is used when all ports are congested. By default, the system enables a maximum of 2 lossless queues on the S4810, S4820T, and MXL platforms, and a maximum of 1 lossless queue on the S6000 platform.

This default behavior is impacted if you modify the total buffer available for PFC or assign static buffer configurations to the individual PFC queues.

Interworking of DCB Map With DCB Buffer Threshold Settings

DCB map functionality is supported on the S4810, S4820T, S6000, I/O Aggregator, and MXL platforms.

The `dcb-input` and `dcb-output` configuration commands are deprecated. You must use the `dcp-map` command to create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic.

You can configure the `dcb-buffer-threshold` command and its related parameters only on ports with either auto configuration or dcb-map configuration. This command is not supported on existing front-panel interfaces or stack ports that are configured with the `dcb-input` or `dcb-output` commands. Similarly, if dcb-buffer-threshold configuration is present on any interface or a stack port, dcb-input or dcb-ouput policy cannot be applied on those interfaces.

Example: When dcb-buffer-threshold is applied on interfaces or stack ports with dcb-input or dcb-output policy, the following error message is displayed:

```
%Error: dcb-buffer-threshold not supported on interfaces with deprecated
commands
```

Example: When dcb-input or dcb-output is configured on interfaces or stack ports with dcb-buffer threshold policy:

```
%Error: Deprecated command is not supported on interfaces with dcb-buffer-
threshold configured
```

You must not modify the service-class dot1p mappings when any buffer-threshold-policy is configured on the system.

```
S4810-1(conf)#service-class dot1p-mapping dot1p0 3
% Error: PFC buffer-threshold policies conflict with dot1p mappings. Please
remove all dcb-buffer-threshold policies to change mappings.
```

The show dcb command has been enhanced to display the following additional buffer-related information:

```
S4810-YU-MR-FTOS (conf)#do show dcb
dcb Status :  Enabled
PFC Queue Count :  2 --Indicate the PFC queue configured.
```

```
 Total buffer (lossy + lossless)(in KB): 7787--Total buffer space for lossy and
lossless queues
 PFC total buffer (in KB): 6526 --Indicates the total buffer (configured or
default)
 PFC shared buffer (in KB): 832--Indicates the shared buffer (Configured or
default)
 PFC available buffer ( in KB): 5694--Indicates remaining available buffers for
PFC that are free to be allocated
```

Configuring the Dynamic Buffer Method

Priority-based flow control using dynamic buffer spaces is supported on the S4810, S4820T, S6000, and MXL platforms.

To configure the dynamic buffer capability, perform the following steps:

1. Enable the DCB application. By default, DCB is enabled and link-level flow control is disabled on all interfaces.

   CONFIGURATION mode

   ```
   S6000-109-FTOS(conf)#dcb enable
   ```

2. Configure the shared PFC buffer size and the total buffer size. A maximum of 4 lossless queues are supported.

   CONFIGURATION mode

   ```
   S6000-109-FTOS(conf)#dcb pfc-shared-buffer-size 4000
   S6000-109-FTOS(conf)#dcb pfc-total-buffer-size 5000
   ```

3. Configure the number of PFC queues.

   CONFIGURATION mode

   ```
   FTOS(conf)#dcb enable pfc-queues 4
   ```

   The number of ports supported based on lossless queues configured will depend on the buffer. The default number of PFC queues in the system is 2 for S4810 and 1 for S6000 platforms.

   For each priority, you can specify the shared buffer threshold limit, the ingress buffer size, buffer limit for pausing the acceptance of packets, and the buffer offset limit for resuming the acceptance of received packets.

4. Configure the profile name for the DCB buffer threshold

   CONFIGURATION mode

   ```
   S6000-109-FTOS(conf)#dcb-buffer-threshold test
   ```

5. DCB-BUFFER-THRESHOLD mode

   ```
   S4810-YU-MR-FTOS(conf-dcb-buffer-thr)# priority 0 buffer-size 52 pause-
   threshold 16 resume-offset 10 shared-threshold-weight 7
   ```

6. Assign the DCB policy to the DCB buffer threshold profile on stack ports.

   CONFIGURATION mode

   ```
   S4810-YU-MR-FTOS(conf)# dcb-policy buffer-threshold stack-unit all stack-
   ports all test
   ```

7. Assign the DCB policy to the DCB buffer threshold profile on interfaces. This setting takes precedence over the default buffer-threshold setting.

   INTERFACE mode (conf-if-te)

   ```
   S4810-YU-MR-FTOS(conf-if-te-0/0)#dcb-policy buffer-threshold test
   ```

8. Create a QoS policy buffer and enter the QoS Policy Buffer Configuration mode to configure the no-drop queues, ingress buffer size, buffer limit for pausing, and buffer offset limit for resuming.

   CONFIGURATION mode

   ```
   S4810-YU-MR-FTOS(conf)# qos-policy-buffer test
   S4810-YU-MR-FTOS (conf-qos-policy-buffer)#queue 0 pause no-drop buffer-size
   128000 pause-threshold 103360 resume-threshold 83520
   S4810-YU-MR-FTOS (conf-qos-policy-buffer)# queue 4 pause no-drop buffer-size
   128000 pause-threshold 103360 resume-threshold 83520
   ```

Applying a DCB Map in a Switch Stack

You must apply the same DCB map with PFC and ETS configuration to all stacked ports in a switch stack. You cannot apply different DCB maps to different stacked switches. This functionality is supported on the S6000 platform.

Entering the **no dcb-map stack-unit all** command removes all PFC and ETS settings applied to stacked ports from the DCB map and resets PFC and ETS to their default settings.

Task	Command	Command Mode
Apply the specified DCB map on all ports of the switch stack.	dcb-map stack-unit all stack-ports all *dcb-map-name*	CONFIGURATION

dcb pfc-shared-buffer-size

Configure the maximum amount of shared buffer size for PFC packets in kilobytes. This utility is supported on the S4810, S4820T, S6000, and MXL platforms.

You must configure the shared buffer size to be less than the total PFC buffer size. If the buffer size and DCB buffer threshold settings are applied on one or more ports, a validaiton is performed to determine whether following condition is satisfied: Shared-pfc-buffer-size <= (Total-pfc-buffer-size - Σpfc priority <> buffer-size on each port, priority).

If the preceding condition is not satisfied by the shared PFC buffer size value, the configuration is not saved and a system logging message is generated stating that the shared buffer size that you attempt to specify cannot be configured because of the existing total buffer space on the system being lower than the shared buffer size. You must either enter a smaller value for the shared buffer size or increase the total buffer size appropriately by using the `dcb pfc-total- buffer-size` command.

S6000 S4810 S4820T MXL

Syntax	`dcb pfc-shared—buffer—size` *KB*	
Parameters	*KB*	Enter a number in the range of 0 to 7787.
Default	The default is 1 KB for S6000 platforms.	
Command Modes	CONFIGURATION mode	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, S6000, and MXL platforms.

Usage Information	Configure the maximum shared buffer available for PFC traffic. You can choose to increase or decrease the shared buffer that is currently allocated in the system by default. You must configure the shared buffer size to be less than the total PFC buffer size. If the buffer size and DCB buffer threshold settings are applied on one or more ports, a validation is performed to determine whether following condition is satisfied:
	
	Shared-pfc-buffer-size <= (Total-pfc-buffer-size - Σpfc priority <> buffer-size on each port, priority).
	
	If the preceding condition is not satisfied by the shared PFC buffer size value, the configuration is not saved and a system logging message is generated as follows:
	
	`S4810-YU-MR-FTOS (conf)#dcb pfc-shared-buffer-size 2000`
	
	`%ERROR: pfc shared buffer size configured cannot accommodate existing buffer requirement in the system.`

Example	`S4810-YU-MR-FTOS (conf)#dcb pfc-shared-buffer-size 5000`

dcb-buffer-threshold

Configure the profile name for the DCB buffer threshold. This utility is supported on the S4810, S4820T, S6000, and MXL platforms.

S6000 S4810 S4820T MXL

Syntax	`dcb buffer—threshold` *`profile-name`*	
Parameters	***profile-name***	Enter the name of the profile, which can be a string of up to 32 characters in length.
Default	None	
Command Modes	CONFIGURATION mode	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, S6000, and MXL platforms.
Usage Information	When you enter the profile name, you enter the DCB buffer threshold configuration mode. You can specify the shared buffer threshold limit, the ingress buffer size, buffer limit for pausing the acceptance of packets, and the buffer offset limit for resuming the acceptance of received packets.	
Example	`S4810-YU-MR-FTOS (conf)#dcb buffer—threshold test`	

priority

Configure the priority for the PFC threshold to be allocated to the buffer space parameters. This utility is supported on the S4810, S4820T, S6000, and MXL platforms.

S6000 S4810 S4820T MXL

Syntax

```
priority value buffer-size size pause-threshold threshold-value
resume-offset threshold-value shared-threshold-weight size
```

Parameters

priority	Specify the priority of the queue for which the buffer space settings apply
value	Enter a number in the range of 0 to 7 to denote the priority to be allocated to the dynamic buffer control mechanism
buffer-size	Ingress buffer size
size	Size of the ingress buffer in KB. Enter a number in the range of 0 to 7787. The default is 45 KB.
pause-threshold	Buffer limit for pause frames to be sent
threshold-value	Buffer limit at which the port sends the pause to peer in KB. Enter a number in the range of 0 to 7787. The default is 10 KB.
resume-offset	Buffer offset limit for resuming in KB
threshold-value	Buffer offset limit at which the port resumes the peer in KB. Enter a number in the range of 1 to 7787. The default is 10 KB.
shared-threshold-weight	Buffer shared threshold weight
size	Weightage of the priorities on the shared buffer size in the system. Enter a number in the range of 0 to 9. The default shared threshold weight is 10.

Default The default size of the ingress buffer is 45 KB. The default buffer limit at which the port sends the pause to peer and recommences the sending of packets to the peer is 10 KB. The default threshold weight of the shared buffer space is 10.

Command Modes DCB-BUFFER-THRESHOLD mode

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, S6000, and MXL platforms.

Usage Information

For each priority, you can specify the shared buffer threshold limit, the ingress buffer size, buffer limit for pausing the acceptance of packets, and the buffer offset limit for resuming the acceptance of received packets. When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device.

You can use the `priority` command to set up both the administrative and peer-related PFC priorities. For example, you can configure the intended buffer configuration for all 8 priorities. If you configure the number of lossless queues as 4 and if the administrator-configured priorities configured within the DCB input policy is applied, then the configuration for those priorities are pre-designed. However, if the peer-provided priorities are applied, although a DCB input policy is present, the peer-provided priorities become effective for buffer configuration. This method of configuration provides an easy and flexible technique to accommodate both administratively-configured and peer-configured priorities.

Example

```
S4810-YU-MR-FTOS (conf-dcb-buffer-thr)#priority 0 buffer-size
52 pause-threshold 16 resume-offset 10 shared-threshold-weight
7
```

qos-policy-buffer

Create a QoS policy buffer and enter the configuration mode to configure the no-drop queues, ingress buffer size, buffer limit for pausing, and buffer offset limit for resuming. This utility is supported on the S4810, S4820T, S6000, and MXL platforms.

S6000 S4810 S4820T MXL

Syntax

```
qos-policy-buffer queue queue-num pause no-drop queue buffer-
size size pause-threshold threshold-value resume-offset
threshold-value shared-threshold-weight size
```

Parameters

policy-name	Name of the QoS policy buffer that is applied to an interface for this setting to be effective in conjunction with the DCB input policy. You can specify the shared buffer threshold limit, the ingress buffer size, buffer limit for pausing the acceptance of packets, and the buffer offset limit for resuming the acceptance of received packets. This method of configuration enables different peer-provided and administrative priorities to be set up because the intended queue is directly configured instead of determining the priority to queue mapping for local and remote parameters.
queue 0 to queue 7	Specify the queue number to which the QoS policy buffer parameters apply
pause	Pause frames to be sent at the specified buffer limit levels and pause packet settings
no-drop	The packets for this queue must not be dropped
value	Enter a number in the range of 0 to 7 to denote the priority to be allocated to the dynamic buffer control mechanism

buffer-size	Ingress buffer size
size	Size of the ingress buffer in KB. Enter a number in the range of 0 to 7787. The default is 45 KB.
pause-threshold	Buffer limit for pause frames to be sent
threshold-value	Buffer limit at which the port sends the pause to peer in KB. Enter a number in the range of 0 to 7787. The default is 10 KB.
resume-offset	Buffer offset limit for resuming in KB
threshold-value	Buffer offset limit at which the port resumes the peer in KB. Enter a number in the range of 1 to 7787. The default is 10 KB.
shared-threshold-weight	Buffer shared threshold weight
size	Weightage of the priorities on the shared buffer size in the system. Enter a number in the range of 0 to 9. The default shared threshold weight is 10.

Default

The default size of the ingress buffer is 45 KB. The default buffer limit at which the port sends the pause to peer and recommences the sending of packets to the peer is 10 KB. The default threshold weight of the shared buffer space is 10.

Command Modes

DCB-BUFFER-THRESHOLD mode

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, S6000, and MXL platforms.

Usage Information

You must apply this buffer policy at the interface level for the attributes to be applicable in conjunction with the DCB input policy.

For each QoS policy buffer, you can specify the shared buffer threshold limit, the ingress buffer size, buffer limit for pausing the acceptance of packets, and the buffer offset limit for resuming the acceptance of received packets. When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device.

You can use set up both the administrative and peer-related PFC priorities. For example, you can configure the intended buffer configuration for all 8 priorities. If you configure the number of lossless queues as 4 and if the administrator-configured priorities configured within the DCB input policy is applied, then the configuration for those priorities are pre-designed. However, if the peer-provided priorities are applied, although a DCB input policy is present, the peer-provided priorities become effective for buffer configuration. This method of configuration provides an easy and flexible technique to accommodate both administratively-configured and peer-configured priorities.

Example

```
S4810-YU-MR-FTOS(conf)# qos-policy-buffer test
S4810-YU-MR-FTOS (conf-qos-policy-buffer)#queue 0 pause no-drop
buffer-size 128000 pause-threshold 103360 resume-threshold
83520
S4810-YU-MR-FTOS (conf-qos-policy-buffer)# queue 4 pause no-
drop buffer-size 128000 pause-threshold 103360 resume-threshold
83520
```

dcb-policy buffer-threshold (Interface Configuration)

Assign the DCB policy to the DCB buffer threshold profile on interfaces. This setting takes precedence over the global buffer-threshold setting. This utility is supported on the S4810, S4820T, S6000, and MXL platforms.

S6000 S4810 S4820T MXL

Syntax

```
dcb-policy buffer-threshold profile-name
```

Parameters

buffer-threshold	Configure the profile name for the DCB buffer threshold
profile-name	Enter the name of the profile, which can be a string of up to 32 characters in length.

Default None

Command Modes INTERFACE mode

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, S6000, and MXL platforms.

Usage Information You can configure up to a maximum of 4 lossless (PFC) queues. By configuring 4 lossless queues, you can configure 4 different priorities and assign a particular priority to each application that your network is used to process. For example, you can assign a higher priority for time-sensitive applications and a lower priority for other services, such as file transfers. You can configure the amount of buffer space to be allocated for each priority and the pause or resume thresholds for the buffer. This method of configuration enables you to effectively manage and administer the behavior of lossless queues.

Example

```
S4810-YU-MR-FTOS(conf-if-te-0/0)#dcb-policy buffer-threshold
test
```

dcb-policy dcb-buffer-threshold (Global Configuration)

Assign the DCB policy to the DCB buffer threshold profile on stack ports that applies globally throughout the system. This utility is supported on the S4810, S4820T, and MXL platforms. This command is not supported on the S6000 platform because it does not contain stack ports.

S4810 S4820T MXL

Syntax

```
dcb-policy buffer-threshold stack-unit all stack-ports all
profile-name
```

Parameters

dcb-buffer-threshold	Configure the profile name for the DCB buffer threshold
profile-name	Enter the name of the profile, which can be a string of up to 32 characters in length.
stack-unit all	Enter the stack unit identification. Indicates the specific the stack unit or units. Entering all shows the status for all stacks.
stack-port all	Enter the port number of a port in a switch stack.

Default None

Command Modes CONFIGURATION mode

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, and MXL platforms.

Usage Information You can configure up to a maximum of 4 lossless (PFC) queues. By configuring 4 lossless queues, you can configure 4 different priorities and assign a particular priority to each application that your network is used to process. For example, you can assign a higher priority for time-sensitive applications and a lower priority for other services, such as file transfers. You can configure the amount of buffer space to be allocated for each priority and the pause or resume thresholds for the buffer. This method of configuraiton enables you to effectively manage and administer the behavior of lossless queues.

Example

```
S4810-YU-MR-FTOS(conf)# dcb-policy buffer-threshold stack-unit
all stack-ports all test
```

show qos dcb-buffer-threshold

Displays the DCB buffer threshold assigned to a QoS policy. This command is supported on the S6000 platform.

Syntax

```
show qos dcb buffer-threshold {name}
```

Parameters

name	Enter the name of the profile, which can be a string of up to 32 characters in length.

Command Modes

EXEC

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S6000 platform.

Usage Information

The following table describes the output fields displayed for the `show` command:

Field	Description
Name	Name of the DCB buffer threshold profile
Buffer threshold parameters	Buffer size allocated for the PFC priority queue and the priority of the queue

Example

```
FTOS#show qos dcb buffer-threshold

Name   :      test1
Buffer threshold parameters:
pfc priority 0 buffer-size 40
pfc priority 3 buffer-size 50
```

show hardware stack-unit buffer-stats-snapshot (With Polling and History)

View the buffer statistics tracking resource information with polling details and historical snapshots. This command is supported on the S6000 platform.

Syntax

```
show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource X history Y
```

Parameters

stack-unit *number*	Unique ID of the stack unit to select a particular stack member and then enter one of the following command options to display a collection of data based on the option entered. The range is 0 to 11.
buffer-stats-snapshot unit *number*	Display the historical snapshot of buffer statistical values
unit	Enter the keyword `unit` along with a port-pipe number, then the keyword counters to display the counters on the selected port-pipe. The range is 0 to 0.
resource *X*	Buffer and traffic manager resources usage, where X can be one of the following: • All - Ingress and Egress resources snapshots

- Port {id |all} queue {all} - egress queue-level snapshot for both unicast and multicast packets
- Port {id |all} queue ucast {id | all} - egress queue-level snapshot for unicast packets only
- Port {id |all} queue mcast {id | all} - egress queue-level snapshot for multicast packets only
- Port {id |all} prio-group {id | all} - ingress priority-group level snapshot

history *Y* — Historical snapshot details of buffer space statistics, where *Y* can be one of the following:

- Instance {all | id} - Displays the information for all instances or the specified instance of the snapshot.
- Summary - Displays the consolidated information pertaining to the preceding three instances of the snapshot values collected in history.

Command Modes

EXEC

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S6000 platform.

Usage Information

When you enter the "instance all" option, "show hardware stack-unit 0 buffer-stats-snapshot unit 0 resource X" Output for all available instances on the history collection is displayed.

When you enter the "instance id" option, "show hardware stack-unit 0 buffer-stats-snapshot unit 0 resource X" for specified instance alone is displayed.

When you enter the "summary" option, "show hardware stack-unit 0 buffer-stats-snapshot unit 0 resource X" will be enhanced to display the total buffered cells, shared cells, headroom cells for last 5 instances in the table format.

If information for specified instance id id is not available when you enter the show command, which occurs if you issue the command before the time elapsed for the snapshot to be captured for that instance ID, the following informational message is displayed on the console:

```
%Info: Data for instance id id is not available.
```

For example, if you configured 5 as the maximum instances with linear periodicity and a polling interval of 10 seconds, 1 as the multiplier, then 5 instances will be polled at 10, 20, 30, 40, and 50 seconds incrementally. If you attempt to enter the show command to display the fifth instance after 30 seconds of enabling polling, the aforementioned information message is shown.

If specified instance ID is higher than the size of the maximum number of snapshot instances configured, the following error message is displayed on the console:

```
%Error: Instance Id is not valid. Configured max snapshot
instances are <max-instances>
```

If you configured the maximum number of instances as 5 and attempt to view the buffer statistics tracking details for the instance ID of 6, the aforementioned error is shown.

Example

```
FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 5 queue all history summary

Stack-unit 0 unit 0 port 5 (interface te 0/4)

-----------------------------------------------------------------
-----------------------------------------------
Q# TYPE     Q#                          TOTAL BUFFERED CELLS
                                  Instance 1  Instance 2
Instance 3 Instance 4 Instance 5
                                  10S            20S
30S           40S           50S
-----------------------------------------------------------------
-----------------------------------------------
UCAST        2          5                4
6                 0              1
UCAST        3          2                0
1                 5              0
UCAST        11         0                3
2                 0              3
MCAST       4          0                0
0                 0              3
```

If only 2 instances are available at the time the above show command is issued, only 2 instances will be displayed in the summary output.

```
-----------------------------------------------------------------
-----------------------------------------------
Q# TYPE     Q#                          TOTAL BUFFERED CELLS
                                  Instance 1  Instance 2
                                  10S            20S
-----------------------------------------------------------------
-----------------------------------------------
UCAST        2          5                4
1
UCAST        3          2                0
UCAST        11         0                3
MCAST       4          0                0

FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 5 prio-group all history summary

Stack-unit 0 unit 0 port 5 (interface te 0/4)

-----------------------------------------------------------------
-----------------------------------------------
PG#    Instance 1      Instance 2      Instance 3
Instance 4      Instance 5
         Shared Hdrm  Shared Hdrm  Shared Hdrm  Shared Hdrm
Shared Hdrm [in CELLS]
-----------------------------------------------------------------
-----------------------------------------------
6      9           2          0          0          1
0          4           1         7
1
```

```
7       0            0          0           0          1
0          0           0         0            0
```

In the following example, the Headroom Cells field indicates the amount of shared buffer area that is allocated to store packets that are received after the pause frame is received or a priority-based flow control pause frame is enabled. When an inbound interface halts the sending of traffic, it must have the buffer space to save all of the packets currently in the buffer, and also all of the packets that were received before the device stops the sending of packets. Headroom space is used for high-priority traffic that needs to be queued and preserved above the input queue limit, such as keepalives and hello messages.

You can use the following sample command output to obtain a consolidated, whole-scale set of statistical counters of buffer resource utilization in the system and identify the ports that you want. All resources will be cleared after their values are displayed.

```
Dell#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource all
Stack-unit: 0 unit: 0 port: 1 (interface Fo 0/0)
---------------------------------------
PG#     SHARED CELLS     HEADROOM CELLS
---------------------------------------
0       0                0
1       0                0
2       0                0
3       0                0
4       0                0
5       0                0
6       0                0
7       0                0

---------------------------------------
Q# TYPE     Q#     TOTAL BUFFERED CELLS
---------------------------------------
UCAST       0      0
UCAST       1      0
UCAST       2      0
UCAST       3      0
UCAST       4      0
UCAST       5      0
UCAST       6      0
UCAST       7      0
UCAST       8      0
UCAST       9      0
UCAST       10     0
UCAST       11     1
MCAST       0      0
MCAST       1      0
MCAST       2      0
MCAST       3      0
MCAST       4      0
MCAST       5      0
MCAST       6      0
MCAST       7      0
MCAST       8      0

Stack-unit: 0 unit: 0 port: 5 (interface Fo 0/4)
---------------------------------------
PG#     SHARED CELLS     HEADROOM CELLS
```

```
----------------------------------------
0       0                0
1       0                0
2       0                0
3       0                0
4       0                0
5       0                0
6       0                0
7       0                0

----------------------------------------
Q# TYPE     Q#     TOTAL BUFFERED CELLS
----------------------------------------
UCAST       0      0
UCAST       1      0
UCAST       2      0
UCAST       3      0
UCAST       4      0
UCAST       5      0
UCAST       6      0
UCAST       7      0
UCAST       8      0
UCAST       9      0
UCAST       10     0
UCAST       11     0
MCAST       0      0
MCAST       1      0
MCAST       2      0
MCAST       3      0
MCAST       4      0
MCAST       5      0
MCAST       6      0
MCAST       7      0
MCAST       8      0
<… snip …>
Stack-unit: 0 unit: 0 port: 104 (interface Te 0/124)
----------------------------------------
PG#     SHARED CELLS     HEADROOM CELLS
----------------------------------------
0       0                0
1       0                0
2       0                0
3       0                0
4       0                0
5       0                0
6       0                0
7       0                0

----------------------------------------
Q# TYPE     Q#     TOTAL BUFFERED CELLS
----------------------------------------
UCAST       0      0
UCAST       1      0
UCAST       2      0
UCAST       3      0
UCAST       4      0
UCAST       5      0
UCAST       6      0
UCAST       7      0
UCAST       8      0
UCAST       9      0
UCAST       10     0
UCAST       11     1
MCAST       0      0
```

```
MCAST       1      0
MCAST       2      0
MCAST       3      0
MCAST       4      0
MCAST       5      0
MCAST       6      0
MCAST       7      0
MCAST       8      0
```

To determine the port that is congested and monitor all queues (including multicast and unicast queues] only on that port:

```
FTOS#$show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 1 queue all
Stack-unit: 0 unit: 0 port: 1 (interface Fo 0/0)
---------------------------------------
Q# TYPE     Q#     TOTAL BUFFERED CELLS
---------------------------------------
UCAST       0      0
UCAST       1      0
UCAST       2      0
UCAST       3      0
UCAST       4      0
UCAST       5      0
UCAST       6      0
UCAST       7      0
UCAST       8      0
UCAST       9      0
UCAST       10     0
UCAST       11     1
MCAST       0      0
MCAST       1      0
MCAST       2      0
MCAST       3      0
MCAST       4      0
MCAST       5      0
MCAST       6      0
MCAST       7      0
MCAST       8      0
FTOS#
```

To examine the port that is congested and monitor all multicast queues on that port:

```
FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 1 queue mcast all
Stack-unit: 0 unit: 0 port: 1 (interface Fo 0/0)
---------------------------------------
Q# TYPE     Q#     TOTAL BUFFERED CELLS
---------------------------------------
MCAST       0      0
MCAST       1      0
MCAST       2      0
MCAST       3      0
MCAST       4      0
MCAST       5      0
MCAST       6      0
MCAST       7      0
MCAST       8      0
FTOS#
```

To determine the port that is congested and monitor all the unicast Queues on that port:

```
FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 1 queue ucast all
Stack-unit: 0 unit: 0 port: 1 (interface Fo 0/0)
---------------------------------------
Q# TYPE     Q#     TOTAL BUFFERED CELLS
---------------------------------------
UCAST       0      0
UCAST       1      0
UCAST       2      0
UCAST       3      0
UCAST       4      0
UCAST       5      0
UCAST       6      0
UCAST       7      0
UCAST       8      0
UCAST       9      0
UCAST       10     0
UCAST       11     0
FTOS#
```

To identify the port that is congested and monitor all the priority groups on that particular port:

```
FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 1 prio all
Stack-unit: 0 unit: 0 port: 1 (interface Fo 0/0)
---------------------------------------
PG#     SHARED CELLS     HEADROOM CELLS
---------------------------------------
0       0                0
1       0                0
2       0                0
3       0                0
4       0                0
5       0                0
6       0                0
7       0                0
FTOS#
```

To determine the specific priority group, unicast or multicast queue that is congested and monitor that queue separately:

```
FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 1 prio 6
Stack-unit: 0 unit: 0 port: 1 (interface Fo 0/0)
---------------------------------------
PG#     SHARED CELLS     HEADROOM CELLS
---------------------------------------
6       0                0
```

dcb pfc-total-buffer-size

Configure the total buffer size for PFC in kilobytes. This utility is supported on the S6000 platform.

S6000

Syntax	`dcb pfc-total—buffer—size KB`
Parameters	***KB*** Enter a number in the range of 0 to 7787.
Default	The default is 1 KB for S6000 platforms.
Command Modes	CONFIGURATION mode
Command History	**Version 9.3.0.0** Introduced on the S6000 platform
Usage Information	Configure the maximum buffer available for PFC traffic. You can choose to increase or decrease the buffer size that is currently allocated in the system by default. However, if you modify the PFC buffer size to be lower than the previously configured PFC buffer size, the system determines whether this reduction in size is valid without disrupting existing configuration. In such a scenario, you must disable and reenable DCB. For example, if you modify the total buffer size to be 4000 KB from the previous size of 5000 KB, an error message is displayed that this reduction cannot be performed owing to existing system configuration because of queues that are being currently processed. The lossless queue limit per port is validated based on the `dcb pfc-queues` command. PFC queue configuration identifies the maximum number of queues a port can support. Although the queue limit per port is a baseline when dynamic buffering is enabled, the limit per port for queues depends on the availability of the buffer.

Example	`S4810-YU-MR-FTOS (conf)#dcb pfc-total-buffer-size 5000`
	
	`S4810-YU-MR-FTOS (conf)#dcb pfc-total-buffer-size 4000 %ERROR: Total pfc buffer size configured cannot accommodate existing buffer requirement in the system.`

show running-config dcb-buffer-threshold

Displays the DCB buffer threshold details in the running configuration. This command is supported on the S6000 platform.

Syntax	`show running—config buffer-threshold`
Command Modes	EXEC EXEC Privilege

Command History

Version 9.3.0.0 Introduced on the S6000 platform.

Usage Information

The following table describes the output fields displayed for the `show running-config dcb-buffer-threshold` command:

Field	Description
Profile name	Name of the DCB buffer threshold profile
Priority	The priority of the queue for which the buffer space settings apply
buffer-size	Ingress buffer size
pause-threshold-value	Buffer limit at which the port sends the pause to peer in KB.
resume-threshold-value	Buffer offset limit at which the port resumes the peer in KB.

Example

```
FTOS#show run buffer-threshold
!
dcb-buffer-threshold test1
pfc priority 0 buffer-size 40
pfc priority 3 buffer-size 50
!
dcb-buffer-threshold test2
pfc priority 0 buffer-size 80 pause-threshold 50
!
dcb-buffer-threshold test3
pfc priority 0 buffer-size 80 pause-threshold 60 resume-
threshold 30


On interface on which PFC is enabled:

Show interface tengigabitethernet 0/0 pfc buffer-threshold
------------------------------------------------------------------
------------------------------
Queue#   Lossless    Buffer-size   Pause-threshold      Resume-
offset        Shared threshold
                     (KB)                (KB)
(KB)          weight
------------------------------------------------------------------
------------------------------
0        No          -                   -
-                    -
1        No          -                   -
-                    -
2        Yes         -                   20
-                    9
3        Yes         52                  25
15                   0
4        Yes         -                   45
25                   5
5        No          -                   -
-                    -
6        No          -                   -
-                    -
7        No          -                   -
```

```
-                    -
-    Denotes dynamic buffering is enabled in respective queues
```

On interface in which PFC is not enabled:

```
FTOS#show interface tengigabitethernet 0/20 pfc buffer-
threshold
```

The following table describes the output fields displayed for the `show interface pfc buffer-threshold` command:

Field	Description
queue	Number of the queue
lossless	Whether the queue is a lossy or lossless queue for which buffer threshold is configured
buffer-size	Ingress buffer size
pause-threshold-value	Buffer limit at which the port sends the pause to peer in KB.
resume-threshold-value	Buffer offset limit at which the port resumes the peer in KB.
shared threshold weight	Weightage of the priorities on the shared buffer size in the system.

dcb pfc-queues

Configure the number of PFC queues. This utility is supported on the S4810 and S6000 platforms.

Syntax `dcb pfc-queues value`

Parameters

value	Enter the number of PFC queues in the range of 0 through 4. The number of ports supported based on lossless queues configured will depend on the buffer.

Default The default number of PFC queues in the system is 2 for S4810 and 1 for S6000 platforms.

Command Modes CONFIGURATION mode

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms

Usage Information You can configure up to a maximum of 4 lossless (PFC) queues. By configuring 4 lossless queues, you can configure 4 different priorities and assign a particular priority to each application that your network is used to process. For example, you can assign a higher priority for time-sensitive applications and a lower priority for other services, such as file transfers. You can configure the amount of buffer space to be allocated for each priority and the pause or resume thresholds for the buffer.

This method of configuration enables you to effectively manage and administer the behavior of lossless queues.

Example

```
Dell(conf)#dcb pfc-queues 4
```

7

Egress Interface Selection (EIS) for HTTP and IGMP Applications

The functionality to configure the egress interface selection (EIS) mechanism is supported on the S4810, S4820T, S6000, and Z9000 platforms.

You can now use EIS to isolate the management and front-end port domains for HTTP and IGMP traffic. Also, EIS enables you to configure the responses to switch-destined traffic with management port IP address as the source IP address to be sent out of the switch through the management port instead of the front-end port.

Management Egress Interface Selection (EIS) feature is applicable only for the out-of-band (OOB) management port. The references to management default route or static route in this chapter denote the routes configured using the management route command. The management default route can be either configured statically or returned dynamically by the DHCP client. A static route points to the Management interface or a forwarding router.

Transit traffic (destination IP not configured in the switch) that is received on the front-end port with destination on the management port is dropped and received in the management port with destination on the front-end port is dropped.

Switch destined traffic (destination IP configured in the switch)

- Received in the front end port with destination IP equal to management port IP address or management port subnet broadcast address is dropped.
- Received in the management port with destination IP not equal to management IP address or management subnet broadcast address is dropped.

Traffic (switch initiated management traffic or responses to switch destined traffic with management port IP address as the source IP address) for user-specified management protocols must exit out of the management port. In this chapter, all the references to traffic indicates switch-initiated traffic and responses to switch-destined traffic with management port IP address as the source IP address.

In customer deployment topologies, it might be required that the traffic for certain management applications needs to exit out of the management port only. You can use EIS to control and the traffic can exit out of any port based on the route lookup in the IP stack.

One typical example is an SSH session to an unknown destination or an SSH connection that is destined to the management port IP address. The management default route can coexist with front-end default routes. If SSH is specified as a management application, SSH links to and from an unknown destination uses the management default route.

Protocol Separation

When you configure the application application-type command to configure a set of management applications with TCP/UDP port numbers to the OS, the following table describes the association between applications and their port numbers.

Table 1. Association Between Applications and Port Numbers

Application Name	Port Number	Client	Server
SSH	22	Supported	Supported
Sflow-Collector	6343	Supported	
SNMP	162 for SNMP Traps (client), 161 for SNMP MIB response (server)	Supported	
NTP	123	Supported	
DNS	53	Supported	
FTP	20/21	Supported	Supported
Syslog	514	Supported	
Telnet	23	Supported	Supported
TFTP	69	Supported	
Radius	1812,1813	Supported	
Tacacs	49	Supported	
HTTP	80 for httpd 443 for secure httpd 8008 http server port for confd application 8888 secure http server port for confd application		Supported

: If you configure a source interface is for any EIS management application, EIS might not coexist with that interface and the behavior is undefined in such a case. You can configure the source interface for the following applications: FTP, ICMP (ping and traceroute utilites), NTP, RADIUS, TACACS, Telnet, TFTP, syslog, and SNMP traps. Out of these applications, EIS can coexist with only syslog and SNMP traps because these applications do not require a response after a packet is sent.

For applications such as RADIUS, TACACS, SSH, and sFlow, user-specified port numbers are also processed by the switch. The OS maintains a list of configured management applications and their port numbers. You can configure two default routes, one configured on the management port and the other on the front-end port.

Two tables, namely, Egress Interface Selection routing table and default routing table, are maintained. In the preceding table, the columns Client and Server indicate that the applications can act as both a client and a server within the switch. The Management Egress Interface Selection table contains all management routes (connected, static and default route). The default routing table contains all management routes (connected, static and default route) and all front-end port routes.

Enabling and Disabling Management Egress Interface Selection

You can enable or disable egress-interface-selection using the management egress-interface-selection command.

When the feature is enabled using the management egress-interface-selection command, the following events are performed:

- The CLI prompt changes to the EIS mode.

- In this mode, you can run the “application” and “no application” commands
- Applications can be configured/unconfigured as management applications using the “application”/ “no application” Cli. All configured applications is considered as management applications and the rest of them as non-management applications.
- All the management routes (connected, static and default) are duplicated and added to the management EIS routing table.
- Any new management route added is installed to both the EIS routing table and default routing table.
- For management applications, route lookup is preferentially done in the management EIS routing table for all traffic. Management port is the preferred egress port. For example, if SSH is a management application, an SSH session to a front-panel port IP on the peer box is initiated via management port only, if the management port is UP and management route is available.
- If SSH request is received on the management port destined to the management port IP address, the response to the request is sent out of the management port by performing a route lookup in the EIS routing table
- If the SSH request is received on the front-end destined to the front-end IP address, the response traffic will be sent by doing route lookup in the default routing table only.
- If the management port is down or route lookup fails in the management EIS routing table, packets are dropped.
- For all non-management applications, traffic exits out of either front-end data port or management port based on route lookup in default routing table.
- Ping and traceroute are always non-management applications and route lookup for these applications is done in the default routing table only.
- For ping and traceroute utilities that are initiated from the switch, if reachability needs to be tested through routes in the management EIS routing table, you must configure ICMP as a management application.
- If ping and traceroute is destined to the Management port IP address, response traffic for these packets are sent by doing route lookup in the EIS routing table.

When the feature is disabled using the no management egress-interface-selection command, the following operations are performed:

- All management application configuration is removed.
- All routes installed in the management EIS routing table are removed.

Handling of Management Route Configuration

When the EIS feature is enabled, the following processing occurs:

- All existing management routes (connected, static and default) are duplicated and added to the management EIS routing table.
- Any management static route newly added using “management route” Cli is installed to both the management EIS routing table and default routing table.
- As per existing behavior, for routes in the default routing table, conflicting front end port routes if configured has higher precedence over management routes. So there can be scenarios where the same management route is present in the EIS routing table but not in the default routing table.
- Routes in the EIS routing table are displayed using the show ip management-eis-route Cli command.

- In the netstat output, the prefix “mgmt” is added to routes in the EIS table so that the user can distinguish between routes in the EIS Routing table and default routing table.
- If the management port IP address is removed, the corresponding connected route is removed from both the EIS routing table and default routing table.
- If a management route is deleted, then the route is removed from both the EIS routing table and default routing table.

Handling of Switch-Initiated Traffic

When the control processor (CP) initiates a control packet, the following processing occurs:

- TCP/UDP port number is extracted from the sockaddr structure in the in_selectsrc call which is called as part of the connect system call or in the ip_output function. If the destination TCP/UDP port number belongs to a configured management application then sin_port of destination sockaddr structure is set to Management EIS ID 2 so that route lookup can be done in the management EIS routing table.
- To ensure that protocol separation is done only for switch initiated traffic where the application acts as client, only the destination TCP/UDP port is compared and not the source TCP/UDP port. Source TCP/UDP port will be a well known port number when the box acts as server
- TFTP is an exception to the above logic.
- For TFTP, data transfer is initiated on port 69, but the data transfer ports are chosen independently by the sender and receiver during initialization of the connection. The ports are chosen at random according to the parameters of the networking stack, typically from the range of temporary ports.
- If route lookup in EIS routing table succeeds, the application specific packet count is incremented. This counter is viewed using the `show management application pkt-cntr` command. This counter is cleared using `clear management application pkt-cntr` command.
- If route lookup in the EIS routing table fails or if management port is down then packets are dropped. The application specific count of the dropped packets is incremented and is viewed using the `show management application pkt-drop-cntr` command. This counter is cleared using `clear management application pkt-drop-cntr` command.
- Packets whose destination TCP/UDP port doesn’t match a configured management application , take the regular route lookup flow in the IP stack.
- In the ARP layer, for all ARP packets received through the management interface, a double route lookup is done, one in the default routing table and another in the management EIS routing table. This is because in the ARP layer we don’t have TCP/UDP port information to decide the table in which route lookup should be done.
- The show arp command is enhanced to show the routing table type for the ARP entry.
- For the clear arp-cache command, upon receving the ARP delete request, the route corresponding to the destination IP is identified. The ARP entries learned in the management EIS routing table are also cleared.
- Therefore, a separate control over clearing the ARP entries learnt via routes in the EIS table is not present. If ARP entry for a destination is cleared in the default routing table, then if an ARP entry for the destionation exists in the EIS table, that entry also will be cleared.
- Because fallback support is removed, if management port is down or route lookup in EIS table fails packets are dropped. Therefore, switch-initiated traffic sessions that used to work previously via fallback may not work now.

Handling of Switch-Destined Traffic

- All traffic received on the management port destined to the management port IP address or received on the front end port destined to the front end IP address is processed by the switch
- If source TCP/UDP port number matches a configured EIS or non-EIS management application and source IP address is Management Port IP address then EIS route lookup is done for the response traffic and hence will be sent out of the management port In this case, source IP address will be management port IP address only if the traffic was originally destined to the management port IP.
- ICMP based applications like ping and traceroute are exceptions to the above logic since we don't have TCP/UDP port number. So if source IP address of the packet matches the management port IP address EIS route lookup is done.
- Management application packet counter is incremented if EIS route lookup succeeds and packet is sent out of the management port.
- If route lookup in the EIS routing table fails or if management port is down then packets are dropped. The management application drop counter is incremented.
- Whenever IP address is assigned to the management port, it is stored in a global variable in the IP stack, which is used for comparison with the source IP address of the packet.
- Rest of the response traffic is handled as per existing behavior by doing route lookup in the default routing table. So if the traffic is destined to the front end port IP address response will be sent out by doing route lookup in the default routing table which is existing behavior.

Consider a sample topology in which ip1 is an address assigned to the management port and ip2 is an address assigned to any of the front panel port.a,b are end users on the management and front panel port networks. The OS-initiated traffic for management applications takes a preference for ip1 as source IP and use the management network to reach the destination. If management port is down or route lookup in EIS routing table fails, ip2 is the source IP and front panel port is used to reach the destination. The fallback route between the management and data networks is used in such a case. At any given time, end-users can access FTOS applications using either ip1 or ip2. Return traffic for such end-user-originated sessions destined to management port ip1 are handled using the EIS route lookup.

Handling of Transit Traffic (Traffic Separation)

This is forwarded traffic where destination IP is not an IP address configured in the switch.

- Packets received on the management port with destination on the front end port is dropped. This is existing behavior.
- Packets received on the front endport with destination on the management port is dropped.
- A separate drop counter is incremented for this case. This counter is viewed using netstat command like all other IP layer counters.

Consider a scenario in which ip1 is an address assigned to the management port and ip2 is an address assigned to any of the front panel port of a switch. End users on the management and front panel port networks are connected. In such an environment, traffic received in the management port destined on the data port network is dropped and traffic received in the front end port destined on the management network is dropped.

Mapping of Management Applications and Traffic Type

The following table summarizes the behavior of applications for various types of traffic when the management egress interface selection feature is enabled.

Table 2. Mapping of Management Applications and Traffic Type

Traffic type / Application type	Switch initiated traffic	Switch destined traffic	Transit Traffic
EIS Management Application	Management is the preferred egress port selected based on route lookup in EIS table . If the management port is down or route lookup fails packets are dropped.	If source TCP/UDP port matches a management application and source IP address is management port IP address, management port is the preferred egress port selected based on route lookup in EIS table . If Management port is down or route lookup fails packets are dropped	Traffic from management port to data port and from data port to management port is blocked
Non-EIS management application	Front-end default route will take higher precedence over management default route and SSH session to an unknown destination uses the front-end default route only. No change in the existing behavior.	If source TCP/UDP port matches a management application and source IP address is management port IP address, management port is the preferred egress port selected based on route lookup in EIS table . If Management port is down or route lookup fails packets are dropped	Traffic from management port to data port and from data port to management port is blocked

- EIS is enabled implies that EIS feature is enabled and the application might or might not be configured as a management application
- EIS is disabled implies that either EIS feature itself is disabled or that the application is not configured as a management application

Transit Traffic

This is the case where traffic is transiting the switch. Traffic has not originated from the switch and is not terminating on the switch.

- Drop the packets that are received on the front end data port with destination on the management port.
- Drop the packets that received on the management port with destination as the front end data port.

Switch-Destined Traffic

This is the case where traffic is terminated on the switch. Traffic has not originated from the switch and is not transiting the switch.

All traffic destined to the switch which is received on management or front end data port is accepted by the switch. Response traffic with Management port IP address as source IP address is handled in the same manner as switch originated traffic.

Switch-Originated Traffic

This is the case where traffic is originating from the switch.

1. Management Applications (Applications that are specifically configured as management applications as defined by this feature):

The management port will be egress port for management applications as defined in this feature. If the management port is down or the destination is not reachable through the management port (Next hop ARP is not resolved etc.) and if the destination is reachable through data port, then the management application traffic is sent out through the front end data port. This is a fallback mechanism that is required.

2. Non Management Applications (Applications that are not configured as management applications as defined by this feature):

   Non-management application traffic will exit out of either front end data port or management port based on routing table. If there is a default route on both the management and front end data port, the default for the data port will be preferred route.

Behavior of Various Applications for Switch-Initiated Traffic

This section describes the different system behaviors that occur when traffic is originating from the switch:

EIS Behavior : If destination TCP/UDP port matches a configured management application route lookup is done in EIS table and management port gets selected as the egress port. If Management port is down or route lookup fails packets are dropped.

EIS Behavior for ICMP : ICMP packets does not have TCP/UDP ports. To do EIS route lookup for ICMP based applications (ping and traceroute), using the source ip option, Management Port IP address should be specified as the source IP address. If Management port is down or route lookup fails packets are dropped.

Default Behavior : Route lookup is done in the default routing table and appropriate egress port is selected.

Protocol	Behavior when EIS is Enabled	Behavior when EIS is Disabled
dns	EIS Behavior	Default Behavior
ftp	EIS Behavior	Default Behavior
ntp	EIS Behavior	Default Behavior
radius	EIS Behavior	Default Behavior
Sflow-collector		Default Behavior
Snmp (SNMP Mib response and SNMP Traps)	EIS Behavior	Default Behavior
ssh	EIS Behavior	Default Behavior
syslog	EIS Behavior	Default Behavior
tacacs	EIS Behavior	Default Behavior
telnet	EIS Behavior	Default Behavior
tftp	EIS Behavior	Default Behavior
icmp (ping and traceroute)	EIS Behavior for ICMP	Default Behavior

Behavior of Various Applications for Switch-Destined Traffic

This section describes the different system behaviors that occur when when traffic is terminated on the switch. Traffic has not originated from the switch and is not transiting the switch. Switch destined traffic is applicable only for applications which acts as server for the TCP session and also for ICMP based applications like ping and traceroute. FTP, SSH, and Telnet are the applications that can function as servers for the TCP session.

EIS Behavior : If source TCP or UDP port matches an EIS management or a non-EIS management application and source IP address is management port IP address, management port is the preferred egress port selected based on route lookup in EIS table . If Management port is down or route lookup fails packets are dropped.

If source TCP/UDP port or source IP address does not match the management port IP address route lookup is done in the default routing table.

EIS behavior for ICMP : ICMP packets does not have TCP/UDP ports. In this case, to perform an EIS route lookup for ICMP based applications (ping and tracerout), you must configure ICMP as a management application. If Management port is down or route lookup fails packets are dropped.

If source IP address does not match the management port IP address route lookup is done in the default routing table.

Default Behavior : Route lookup is done in the default routing table and appropriate egress port is selected.

Protocol	Behavior when EIS is Enabled	Behavior when EIS is Disabled
ftp	EIS Behavior	Default Behavior
http	EIS Behavior	Default Behavior
ssh	EIS Behavior	Default Behavior
Snmp (snmp mib response)	EIS Behavior	Default Behavior
telnet	EIS Behavior	Default Behavior
icmp (ping and traceroute)	EIS Behavior for ICMP	Default Behavior

Interworking of EIS With Various Applications

Stacking

- The management EIS is enabled on the master and the standby unit.
- As Traffic can be initiatedfrom the Master unit only, the preference to management EIS table for switch initiated traffic and all its related ARP processing is done in the Master unit only.
- ARP related processing for switch destined traffic is done by both master and standby units.

VLT

VLT feature is for the front end port only. As this feature is specific to the management port, this feature can co-exist with VLT and nothing specific needs to be done in this feature to handle VLT scenario.

DHCP

- If DHCP Client is enabled on the management port, a management default route is installed to the switch.

- If management EIS is enabled, this default route is added to the management EIS routing table and the default routing table.

ARP learn enable

- When ARP learn enable is enabled the switch shall learn ARP entries for ARP Request packets even if the packet is not destined to an IP configured in the box.
- ARP learn enable feature shallnot be applicable to the EIS routing table. It is applicable to the default routing table only. This is to avoid unnecessary double ARP entries

Sflow

Sflow management application is supported only in standalone boxes and switch shall throw error message if sflow is configured in stacking environment

application (for HTTP and ICMP)

Configure the management egress interface selection for HTTP and ICMP.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective Command Reference Guide of the applicable platform of the Release 9.2.0.0 documentation set.

Z9000 S4810 S4820T

Syntax

`application {all | application-type}`

To remove a management application configuration, use the `no application {all | application-type}` command.

Parameters

application-type	Enter any of the following keywords: • For HTTP, enter the keyword `http`. • For ICMP, enter the keyword `icmp`.
all	Configure all applications.

Defaults

None.

Command Modes

EIS Mode (conf-mgmt-eis)

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.3.(0.0)	Added support for the HTTP and ICMP traffic on the Z9000, S4810, and S4820T.

8

Flex Hash and Optimized Boot-Up

This chapter describes the Flex Hash and fast-boot enhancements and contains the following sections:

- Optimizing the Boot Time
- Flex Hash Capability Overview
- Preserving 802.1Q VLAN Tag Value for Lite Subinterfaces

Flex Hash Capability Overview

This functionality is supported on the S6000 platform.

With the introduction of various overlay technologies such as network virtualization using generic routing encapsulation (NVGRE) segments and Routable Remote Direct Memory Access (RRDMA) over Converged Ethernet (RRoCE), information related to a traffic flow is contained in the L4 header. The fields in the L2 and L3 headers are not sufficient to distinguish the flows. Therefore, the fields in the L4 header are processed when hashing is performed for packets over LAG and ECMP links. The Flex Hash functionality enables you to configure a packet search key and matches packets based on the search key. When a packet matches the search key, two 16-bit hash fields are extracted from the start of the L4 header and provided as inputs (bins 2 and 3) for RTAG7 hash computation. You must specify the offset of hash fields from the start of the L4 header, which contains a flow identification field.

You can cause the system to include the fields present at the offsets that you define (from the start of the L4 header) as a part of LAG and ECMP computation. Also, you can specify whether the IPv4 or IPv6 packets must be operated with the Flex Hash mechanism.

Keep the following points in mind when you configure the Flex Hash capability:

- A maximum of eight flex hash entries is supported.
- A maximum of 4 bytes can be extracted from the start of the L4 header.
- The offset range is 0 – 30 bytes from the start of the L4 header.
- Flex Hash uses the RTAG7 bins 2 and 3 (overlay bins). These bins must be enabled for Flex Hash to be configured. These bins contain the source module and source port information. These bins are disabled by default in releases of Dell Networking OS earlier than Release 9.3.0.0. The default behavior of disabling of these bins occurs because of incorrect egress port information that would otherwise be displayed in the output of the diagnostic show command of show ip flow.
- If you configure the Flex Hash mechanism by using the `load-balance ingress-port enable` and the `load-balance flexhash` commands, the `show ip flow` and `show port-channel-flow` commands are not operational. Flex Hash settings and these show commands that display the Layer 3 packets and Layer 2 packets forwarding and flows are mutually exclusive; only either of these capabilities can be functional at a point in time.

  This behavior occurs because the Flex Hash capability is disabled by default, which causes the proper functioning of `show ip flow` and `show port-channel-flow` commands. For the Flex Hash algorithm to work on the S6000 platform, you must enter the `load-balance ingress-port enable` command, which preempts the usability of the IP or Layer 2 trace flow functionalities. This

condition occurs owing to hardware limitations in the S6000 platform, in which the RTAG7 hash selection bitmap overlay bits 2 and 3 need to be enabled for the Flex Hash algorithm and to be disabled for IP and Layer 2 trace flow feature. IP and Layer 2 trace flow feature is useful in identifying the egress interface that the packet uses to pass through or traverse for port-channel and ECMP links. If these overlay bits are enabled, the hashing algorithm calculation contains the source module and source port ID, which causes an incorrect hash value to be computed for the flow packets.

load-balance ingress-port enable

Enable the Flex hash functionality. This utility is supported on the S6000 platform.

Syntax	`load-balance ingress-port enable` To disable the Flex hash capability, use the **no** version of this command.
Default	None
Command Modes	CONFIGURATION mode
Command History	**Version 9.3.0.0** Introduced on the S6000 platform
Usage Information	Flex hash uses the RTAG7 bins 2 and 3 (overlay bins). These bins must be enabled for Flex hash to be configured. These bins contain the source module and source port information. These bins are disabled by default in releases of Dell Networking OS earlier than Release 9.3.0.0. The default behavior of disabling of these bins occurs because of incorrect egress port information that would otherwise be displayed in the output of the diagnostic show command of `show ip flow`. As a result, when load balancing of RRoCE packets using Flex hash is enabled, the `show ip flow` command is not functional. Similarly, when `show ip flow` command operates (ingress port based load balancing is disabled) the hashing of RRoCE packets is not operational. Flex hash APIs do not mask out unwanted byte values after extraction of the data from the Layer 4 headers for the offset value.
Example	`S4810-YU-MR-FTOS(conf)#load-balance ingress-port enable`

load-balance flexhash

Specify the parameters for the Flex Hash mechanism, such as whether IPv4 or IPv6 packets must be subject to Flex Hash functionality, a unique protocol number, the offset of hash fields from the start of the L4 header to be used for hash calculation, and a meaningful description to associate the protocol number with the name. This utility is supported on the S6000 platform.

Syntax	`load-balance flexhash ipv4/ipv6 ip-proto <protocol number> <description string> offset1 <offset1 value> [offset2 <offset2 value>]` To disable the Flex hash settings, use the `no load-balance flexhash ipv4/ipv6 ip-proto protocol number` command.

Parameters

ipv4	Denotes whether Flex Hash needs to be enabled for IPv4 packets.
ipv6	Denotes whether Flex Hash needs to be enabled for IPv6 packets.
protocol number	Represents the Outer IPv4 protocol field in case of IPv4 packets, and the Outer IPv6 next header field in case of IPv6 packets. The `ipv4/ipv6` keyword and the IP protocol value are used as keys to identify if a duplicate flex hash configuration is already present. Duplicate flex hash configuration is not possible. To change an existing flex hash configuration, you must delete the existing flex hash attribute and configure the flex attribute afresh.
description string	A description string is followed by the protocol number to enable you to associate the protocol number with the protocol name in an easily-identifiable way. For example, for a protocol number of 254, you can specify the description as RRoCE.
offset1 *value*	Specify the byte offset from the start of the L4 header from which the 2-byte data is extracted and be used in hash computation. You must enter the offset as an even number. The offset range is 0 – 30 bytes from start of L4 header.
offset2 *value*	(Optional) Specify the additional 2 bytes that must be extracted from the start of the L4 header to be used for hash computation. You must enter the offset as an even number. The offset range is 0 – 30 bytes from start of L4 header.

Default None

Command Modes CONFIGURATION mode

Command History

Version 9.3.0.0	Introduced on the S6000 platform

Usage Information

With the introduction of various overlay technologies such as network virtualization using generic routing encapsulation (NVGRE) segments and Routable Remote Direct Memory Access (RRDMA) over Converged Ethernet (RRoCE), information related to a traffic flow is contained in the L4 header. The fields in the L2 and L3 headers are not sufficient to distinguish the flows. Therefore, the fields in the L4 header are processed when hashing is performed for packets over LAG and ECMP links. The Flex Hash functionality enables you to configure a packet search key and matches packets based on the search key. When a packet matches the search key, two 16-bit hash fields are extracted from the start of the L4 header and provided as inputs (bins 2 and 3) for RTAG7 hash computation. You must specify the offset of hash fields from the start of the L4 header, which contains a flow identification field.

You can cause the system to include the fields present at the offsets that you define (from the start of the L4 header) as a part of LAG and ECMP computation. Also, you can specify whether the IPv4 or IPv6 packets must be operated with the Flex Hash mechanism.

Example

```
S4810-YU-MR-FTOS(conf)# load-balance flexhash ipv4 ip-proto 1
desc offset1 1 offset2 2
```

Configuring the Flex Hash Mechanism

This configuration is supported on the S6000 platform.

The Flex Hash functionality enables you to configure a packet search key and matches packets based on the search key. When a packet matches the search key, two 16-bit hash fields are extracted from the start of the L4 header and provided as inputs (bins 2 and 3) for RTAG7 hash computation. You must specify the offset of hash fields from the start of the L4 header, which contains a flow identification field.

1. In Dell Networking OS Release 9.3.0.0, you can enable bins 2 and 3 by using the load-balance ingress-port enable command in Global Configuration mode. To configure the Flex hash functionality, you must enable these bins.

   CONFIGURATION mode

   ```
   S6000-109-FTOS(conf)# load-balance ingress-port enable
   ```

   As a result, when load balancing of RRoCE packets using Flex hash is enabled, the `show ip flow` command is not functional. Similarly, when `show ip flow` command operates (ingress port based load balancing is disabled) the hashing of RRoCE packets is not operational.

   Flex hash APIs do not mask out unwanted byte values after extraction of the data from the Layer 4 headers for the offset value.

2. You can use the `load-balance flexhash` command to specify whether IPv4 or IPv6 packets must be subject to Flex Hash functionality, a unique protocol number, the offset of hash fields from the start of the L4 header to be used for hash calculation, and a meaningful description to associate the protocol number with the name.

   CONFIGURATION mode

   ```
   Dell(conf)# load-balance flexhash ipv4/ipv6 ip-proto <protocol number>
   <description string> offset1 <offset1 value> [offset2 <offset2 value>]
   ```

   To delete the configured flex hash setting, use the no version of the command.

Configuring Fast Boot and LACP Fast Switchover

To configure the optimized booting time functionality, and quicker aggregation and convergence of member ports of a port- channel bundle, perform the following steps. This procedure is supported on the S6000 platform.

1. Enable the system to be restarted during the next reboot of the device with optimized booting-time functionality enabled. When you restart the device in fast boot mode, traffic disruption is reduced significantly and the system operations to service the data traffic are restored in a seamless way.

   CONFIGURATION mode

   ```
   Dell(conf)#reload-type fastboot
   ```

2. Cause the physical ports to be aggregated faster by configuring this capability in a port-channel on both the nodes that are members of a port-channel. You can configure the optimal switchover functionality for LACP even if you do not enable the fast boot mode on the system. This command applies to dynamic port-channel interfaces only. When applied on a static port-channel, this command has no effect. If you configure the optimized booting-time capability and perform a reload of the system, the LACP application sends PDUs across all the active LACP links immediately.

   INTERFACE (conf-if-po-number) mode

   ```
   Dell(conf-if-po-number)#lacp fast-switchover
   ```

reload-type fastboot

Restart the system with optimized booting-time functionality enabled. When you restart the device in fast boot mode, traffic disruption is reduced significantly and the system operations to service the data traffic are restored in a seamless way. This command is supported on the S6000 platform.

S6000

Syntax	`reload-type fastboot`	
Parameters	**fastboot**	Enable the system to restart the next time with the optimized booting-time capability
Defaults	By default, the device reloads in Jumpstart or BMP mode.	
Command Modes	GLOBAL CONFIGURATION	
Command History	**Version 9.3.0.0**	Added support for the `fastboot` parameter for S6000 platform.
Usage Information	You can configure an optimization technique to reduce the booting time of an S6000 Switch. This mechanism is also called fast boot. With the reduced time that is taken to reboot the switch, upon a manually-initiated reload or an expected restart of the device, the disruption in traffic that is serviced by the switch is minimized. Traffic outage is lowered considerably (reduced to approximately 25 seconds in certain network deployments) when you enable this optimization method for booting of the device. By reducing the duration of traffic loss, subscriber sessions are processed and preserved in an effective and seamless way.	

Related Commands	• show reload-type — displays the current reload mode (BMP or Normal mode). • stop bmp — stops the BMP process and prevents a loop if the DHCP server is not found.

lacp fast-switchover

Cause the physical ports to be aggregated faster by configuring this capability in a port-channel on both the nodes that are members of a port-channel.

S6000

Syntax	`lacp fast-switchover` To disable the capability of faster aggregation of the member ports of a LAG or a port-channel bundle, use the no version of this command.
Defaults	Not configured
Command Modes	INTERFACE (conf-if-po-number)
Command History	**Version 9.3.0.0** Introduced on the S6000.
Usage Information	You can configure the optimal switchover functionality for LACP even if you do not enable the fast boot mode on the system. You must configure the long timeout mechanism for the LACP session to enable the fast boot capability to operate properly. This command applies to dynamic port-channel interfaces only. When applied on a static port-channel, this command has no effect. If you configure the optimized booting-time capability and perform a reload of the system, the LACP application sends PDUs across all the active LACP links immediately.
Related Commands	show lacp — displays the LACP configuration.

Optimizing the Boot Time

This functionality is supported on the S6000 platform.

You can configure an optimization technique to reduce the booting time of an S6000 Switch. This mechanism is also called fast boot. With the reduced time that is taken to reboot the switch, upon a manually-initiated reload or an expected restart of the device, the disruption in traffic that is serviced by the switch is minimized. Traffic outage is lowered considerably (reduced to approximately 25 seconds in certain network deployments) when you enable this optimization method for booting of the device. By reducing the duration of traffic loss, subscriber sessions are processed and preserved in an effective and seamless way.

You can configure this capability on an S6000 Switch that is deployed as a top-of-rack (ToR) switch. The ToR switch is the single point of connection to the network for servers in that rack. This functionality of minimized reload time is supported in a network deployment in which the servers are connected through a ToR, leaf and spine unit or configuration setup. An exterior border gateway protocol (EBGP) session

exists between the ToR and leaf switch units, and between the leaf and spine units or nodes. For example, you can enable the optimized booting method in a deployment in which Micrsoft Bing and Microsoft Azure applications are installed, although different QoS configurations might be needed because of Bing being a search utility and Azure being a service provider that hosts a public cloud.

Consider a sample scenario in which to the south or lower end of the ToR switch, which is an S6000 Switch, the storage servers are connected. To the north or the upper end of the ToR switch, leaf nodes are connected. Spine nodes are at the top of the vertical hierarchical configuration and are connected to the leaf nodes. You can connect up to 96 physical servers in 4 to 5 different subnets, and up to 8 Multiprocotol BGP (MP-BGP) sessions to the servers, that function as load balancers. All the servers are single-homed servers, which does not provide redundancy to servers if a ToR switch fails. The servers advertise both IPv4 and IPv6 addresses. Layer 2 network is to the south of ToR and Layer 3 network is to the north of ToR. An EBGP hold timer of 10 seconds and BGP graceful restart are specified. A maximum of 4000 routes each for Ipv4 and IPv6 traffic can be supported.

To the north of the ToR switch, up to 8 leaf nodes are connected. Different EBGP sessions for IPv4 and IPv6 for each leaf node are configured. LACP is enabled between the ToR and leaf nodes, and the LACP long timer is set to the default value. You can enable the optimized boot functionality in such a topology.

Booting Process When Optimized Boot Time Mechanism is Enabled

When a S6000 switch running Dell Networking OS earlier than Release 9.3.0.0 is reloaded, the CPU and other components on the board are reset at the same time. Therefore, the control plane and the forwarding plane are impacted immediately. After the system boots up and reinitializes, the interfaces come up, control plane protocols are reestablished, network topology information (such as routes, adjacency settings) is learned and installed before the traffic resumes. It is observed that in a typical network scenario, a traffic disconnection of 150 seconds or more occurs. When you employ the optimized booting functionality, the traffic outage duration is reduced in a sizeable way.

Guidelines for Configuring Optimized Booting Mechanism

Keep the following points and limitations in mind when you configure the capability to minimize the booting time:

- The Fastboot functionality is supported only when you perform an expected, stipulated reload by using the reload-type normal-reload command in Global Configuration mode or by using the reset command in uBoot mode on a switch that is running Dell Networking OS Release 9.3.0 or when you perform a planned upgrade (and not an abrupt or unexpected shutdown) from an older release of Dell Networking OS to Release 9.3.0.0 or later, which supports the fast boot capability. We recommend that you do not perform a downgrade of your system from Release 9.3.0.0 to an earlier release that does not support the fast boot functionality. If you downgrade the system to a release earlier than Release 9.3.0.0, the system behavior is unexpected and undefined.
- The Fastboot functionality uses the Symmetric Multiprocessing (SMP) utility that is enabled on the Intel CPU on the S6000 Switch to enhance the speed of the system startup. SMP is supported on the S6000 platform.

For the fastboot mechanism to reduce the traffic disruption significantly, the following conditions apply:

1. When LACP is used between the ToR switch and the adjacent devices, LACP is configured on these adjacent devices with a timeout value of 90 seconds or longer.
2. BGP timers between the ToR switch and adjacent devices are set to high values (for example, a hold timeout of 180 seconds) unless BGP graceful Restart is used.

3. Before performing the planned reload, we recommend that the IPv6 Neighbor Discovery (ND) reachable timer (amount of time that a switch can connect to a remote node after a reachability confirmation event has taken place) is increased to a value of 300 seconds or longer on the adjacent devices to prevent the ND cache entries from becoming stale and being removed while the ToR goes through a CPU reset. This timer can be restored to its prior value after the ToR has completed its planned reload.
4. BGP protocol on the adjacent devices responds to network (link-state) changes and route advertisements quickly and propagate these further up the network quickly. You might need to adjust the BGP timers on these devices.
5. Note that fastboot will operate even if some of the preceding conditions are not met. However, the duration of traffic loss might be longer.
6. Warm boot is supported because it enables faster convergence and reduced traffic loss.
7. BGP graceful restart must be configured with GR time left to default (120 seconds) or higher. The BGP hold timer to be configured with 10 seconds.
8. You must configure the LACP long timeout, which is the amount of time that a LAG interface waits for a PDU from the remote system before bringing the LACP session down, to be higher than the default value.
9. Traffic from North-South and South-North nodes are of line rate type.
10. Traffic outage for a planned reboot is less than 30 seconds for 4000 routes of IPv4 and IPv6 traffic for all of the following traffic directions. These traffic patterns apply only to the S6000 platforms.

- South-North
- North-South
- East-West
- West-East

To the south of ToR, 96 servers can be linked. Up to 8 Multiprocotol BGP (MP-BGP) sessions to the servers are established. You can configure a minimum of 2 MP-BGP sessions and a maximum of 8 MP-BGP sessions.

To the north of the ToR switch, up to 8 leaf nodes are connected. Up to 8 EBGP sessions for IPv4 and IPv6 for each leaf node are configured. LACP is enabled between the ToR and leaf nodes, and the LACP long timer is set to the default value. You must configure 96 ports to be 10-Gigabit Ethernet interfaces and 8 ports as 40-Gigabit ethernet interfaces. You must configure the switch to operate with an uplink speed of 40 Gigabit Ethernet per second.

Interoperation of Applications with Fast Boot and System States

This functionality is supported on the S6000 platform.

The S6000 switch contains a Complex Programmable Logic Device (CPLD) that supports individual device resets in addition to the ability to reset the entire board. This is the underlying, principal capability that the Fastboot functionality is built on. To enable this capability in a usert-initiated manner, you can use the fastboot keyword with the reload-type command in Global Configuration mode in Dell Networking OS Release 9.3.0.0 and later. If you enable the optimzied booting or fast boot functionality, only the CPU is reset. The BIOS is first run, which causes the boot loader (GRUB) located in the Compact Flash card to be started. GRUB loads and runs the appropriate OS image as specified in the boot configuration. The switch will be reset and reinitialized only in the early stage of initialization of this OS

image, after the kernel initialization is complete. Therefore, the forwarding plane match occurs at a later time than the stage at which the match occurs when fast boot capability is not configured.

The following sections describe the working behavior of applications when fast boot functionality is enabled and the various system conditions:

LACP and IPv4 Routing

The following events occur when the operator initiates a fastboot (prior to the CPU being reset) when IPv4 routing is enabled on the system:

The system saves all dynamic ARP entries to a database on the flash drive.

A file is generated to indicate that the system is undergoing a fast boot, which is used after the system comes up.

After the OS FTOS image has been loaded and activated, and the appropriate software components come up, the following additional actions are performed:

- If a database of dynamic ARP entries is present on the flash drive, that information is read and the ARP entries restored ; the entries are installed on the switch as soon as possible. At the same time, the entries are changed to an initial ("aged out") state so that they are refreshed (and flushed if not learnt again). The database on the flash card is also deleted instantaneously.
- The system ensures that local routes known to BGP (configured through the `network or redistribute` commands) are imported into BGP quickly and advertised to peers as quickly as possible. In this process, any advertisement-interval configuration is not considered (only during the initial period when the peer comes up).

If you do not configure BGP GR, you must configure the peering with BGP keepalive and hold timers to be as high as possible (depending on your network deployment and the scaled parameters or sessions) to enable the connection to be active until the ToR reinitializes the switch causing the links to adjacent devices to go down. If the BGP sessions are disabled before the reinitialization of the switch occurs owing to the timeout of the peer, traffic disruption occurs from that point onwards, although the ToR continues to maintain valid routing information in hardware and is capable of forwarding traffic.

LACP and IPv6 Routing

The operation of the fast boot mechanism when the system is configured with IPv6 interfaces and has IPv6 routes is similar to the processing that is done with IPv4 routing and LACP configured. The following IPv6-related actions are performed during the reload phase:

- The system saves all of the dynamic ND cache entries to a database on the flash card. After the system comes back online and the OS image is loaded and the corresponding software applications on the system are also activated, the following processes specific to IPv6 are performed:
- If a database of dynamic ND entries is present on the flash, the same information is read and the ND entries restored (to the IPv6 subsystem as well as the kernel); the entries are installed on the switch as quickly as possible. At the same time, the entries are changed to an initial ("incomplete") state so that they are refreshed (and flushed if not learnt again). The database on the flash is also deleted immediately.
- To ensure that the adjacent systems do not time out and purge their ND cache entries, the age-out time or the reachable time for ND cache entries must be configured to be as high as necessary. We recommend that you configure the reachable timer to be 90 seconds or longer.

BGP Graceful Restart

The fast boot functionality operates in the following manner when the system contains one or more BGP peerings configured for BGP graceful restart, apart from performing the other generic system-wide tasks:

When you reload the device using the fast boot capability enables, a closure of the TCP sessions is performed on all sockets corresponding to BGP sessions on which Graceful Restart has been negotiated. This behavior is to force the peer to perform the helper role so that any routes advertised by the restarting system are retained and the peering session will not go down due to BGP Hold timeout (which needs to be prevented because it does not cause graceful restart to be initiated).

Termination of TCP connections is not initiated on BGP sessions without GR because such a closure might cause the peer to immediately purge routes learnt from the restarting ToR, thereby causing immediate traffic loss

When BGP on the system is started, if it determines that the system has come up through a fastboot, it sets the R-bit and F-bit in the GR capability when bringing up the session with peers for which BGP GR has been configured. This is the standard behavior of a Restarting system and ensures that the peer continues to retain the routes previously advertised by the system.

The system will also delay sending the BGP End-of-RIB notification to peers with whom BGP GR has been negotiated to ensure that the local routes of the system are advertised to the peers, if required by the configuration. This condition occurs only if the system has come up through a fastboot.

Note that if BGP GR is enabled on any peering session, the timeout values used for the BGP hold timer do not take effect.

Cold Boot Caused by Power Cycling the System

When you perform a power-cycle operation on a system that is configured with the optimized booting functionality, in a sequenced, planned manner by powering off the system and then turning it back on and booting it, the system will go through its regular boot sequence as described in the Booting Process When Optimized Boot Time Mechanism is Enabled' section occurs even if it is configured for fastboot. When the system comes up, it is expected that there will be no dynamic ARP or ND database to restore. Likewise, the mode of boot up of the system will not be fastboot and actions specific to this mode (listed in earlier sections) will not be performed.

Unexpected Reload of the System

When an unexpected or unplanned reload occurs, such as a reset caused by the software, the system performs the regular boot sequence ás described in the 'Booting Process When Optimized Boot Time Mechanism is Enabled' section even if it is configured for fastboot. When the system comes up, dynamic ARP or ND database entries are not present or required to be restored. Also, the boot mode of the system is not of fast boot type and the processes that are performed during a normal reload of the system, without fast boot capability enabled, are run.

Software Upgrade

The system behavior when fast boot is used to upgrade the system to a Dell Networking OS release that supports the fastboot functionality enables the restoration of dynamic ARP or ND databases that were maintained in the older release from when you performed the upgrade and the ARP and ND applications identify that the system has been booted using the fast boot funcionality.

LACP Fast Switchover

For the fastboot functionality, the operation of LACP has been optimized. These LACP optimizations are applicable even when fast boot mechanism is not activated when a system reload is performed. These enhancements are controlled using the fast-switchover option that is available with the lacp command in Port Channel Interface Configuration mode. When LACP 'fast-switchover' is enabled on the system, two optimizations to LACP behavior of the local system are performed:

- The wait-while timer is not started in the 'waiting' state of the MUX state machine. Instead the port moves directly to the 'attached' state.
- The local system moves to the 'collecting' and 'distributing' states on the port in a single step without waiting for the partner to set the 'collecting' bit.

The aforementioned optimizations do not work individually to reduce the time that is required for LACP to consider a port as aggregated to a port channel because the partner (adjacent system) is also involved in this process to minimize the switchover time when a failover occurs from one member interface of the LAG or port channel bundle to another member interface.

Changes to BGP Multipath

When the ToR switch becomes active after a system restart using the fast boot method, a change has been made to the BGP multipath and ECMP behavior. The ToR delays the computation and installation of additional paths to a destination into the BGP routing information base (RIB) and forwarding table for a certain period of time. This method of processing occurs to ensure that the ToR is able to learn and install at least one path to each destination as quickly as possible into the forwarding table and to avoid delays that might potentially occur in resuming traffic to some destinations because multiple paths are being computed and installed to other destinations. This operation also occurs to handle a few behavior implications of the FIB application. Additional paths will be automatically computed and installed (if any) without the need for any manual intervention after 30 seconds of the system returning online after a restart, after all established peers have synchronized with the restarting ToR, or a combination of both the conditions.

One possible impact of this behavior change is that if the amount of traffic to a destination is higher than the volume of traffic that can be carried over one path, a portion of that traffic might be dropped for a short duration (30-60 seconds) after the ToR switch comes up. However, this brief traffic disruption is an effective alternative to a complete or increased disconnection in traffic to some destinations for a longer period after the ToR switch returns to the active state.

Minimized Connection Setup Time

Connection establishment is enhanced by performing a retry every second instead of performing a retry to set up the connection at an interval between 15 and 20 seconds with a backoff timer. A faster retry occurs only if the system comes online using the fast boot functionality and only for the initial connection setup (first establishment for each peer). Also, this phenomenon occurs only for a maximum of 60 retries (approximately 60 seconds with a retry every second). If the peering session is not established, the behavior for setting up of the connection is the same as the behavior without fast boot functionality configured.

Faster Local Route Aadvertisements

Local routes from the routing table (RTM) can be injected into BGP using either "redistribute" or "network" configuration. For remote traffic to reach local destinations faster, these routes need to be injected into

BGP as quickly as possible and then advertised to peers as soon as the peering is established. This rapidness in the transmission of routes is essential when BGP GR is not used because the peers will not have the routes from the restarting ToR. The following design modifications have been performed:

Poll for routes corresponding to "networks" every 3 seconds for the first 20 seconds after BGP starts. After that, revert to the usual FTOS behavior of checking every 30 seconds. When doing faster polling, only poll for networks that are not already injected.

Process redistributed routes from RTM every 1 second for the first 20 seconds after BGP starts. After that period, processing of routes occurs every 10 seconds.

To ensure that the local routes are advertised to peers quickly, any configured (or default) neighbor advertisement interval will not be enforced for the first 60 seconds after a peering is established.

Delayed Installation of ECMP Routes Into BGP

The current FIB component of FTOS has some inherent inefficiency when handling a large number of ECMP routes (i.e., routes with multiple equal-cost next hops). To circumvent this, for the specific target configuration for Fastboot, changes are made in BGP to delay the install of ECMP routes. This is done only if the system comes up through a fastboot reload. The BGP route selection algorithm only selects one best path to each destination and delays install of additional ECMP paths until a minimum of 30s from the time the first BGP peer is established. Once this time has elapsed, all routes in the BGP RIB are processed for additional paths.

While the above change will ensure that at least one path to each destination gets into the FIB as quickly as possible, it does prevent additional paths from being used even if they are available. This downside has been deemed to be acceptable.

Changes for BGP Graceful Restart Processes

If BGP Graceful Restart is enabled between the restarting ToR and any adjacent peer, the ToR can take advantage of it and trigger the peer to retain the routes advertised by the ToR even if it goes down. The following enhancements have been made:

1. To trigger the peer router to go into the role of a GR Helper router, the BGP code registers for and receives a notification of operator-initiated fastboot reload. Upon getting this notification, BGP will initiate a TCP close towards every neighbor with whom GR has been negotiated which should cause the peer to transition to the role of a GR Helper. Note that without this change, the peer would detect a change either upon a Hold timeout or when the link to the ToR goes down and neither event would trigger any BGP GR actions on the peer.
2. Upon BGP starting up, if it determines that the system has come up through fastboot, the R-bit and F-bit (Restarting bit and Forwarding bit) are set in the GR capability exchanged with GR-enabled peers. This behavior notifies the peer that the restarting ToR has preserved the forwarding state and ensures that the peer continues to retain the routes previously advertised by the ToR.
3. For all peers that have been established with GR-enabled after a fastboot, the End-of-RIB is delayed by 60 seconds (from the time of establishment of the first peer). This is done to ensure that the ToR would have advertised all its local routes to the peer before sending the EoR.

Operation of LACP

A set of optimizations has been implemented in LACP to speed up the aggregation of ports. However, for these changes to produce a material effect, the peer router also has to behave in an expedient manner (for example, support the same changes). Therefore, these changes may result in maximum benefits when all switches involved are Dell and run the Yakima release.

1. The ‘wait-while’ timer is not started in the ‘waiting’ state of the MUX state machine. The standard recommends waiting for some time for additional ports to try to join the aggregator as that may potentially cause the original port to be unselected or be placed in a ‘standby’ state. However, FTOS does not support the concept of a ‘standby’ state and all ports that are operational can be attached to the aggregator. Therefore, this timer is not started and instead the port moves directly to the ‘attached’ state.
2. Optimization to the MUX state machine provides coupled control instead of independent control. Because of this phenomenon, the state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution. This process optimizes a PDU exchange.

Operation of FIB

For the FIB application, the CAM index generation for NextHops and FirstHops is enhanced to use a bitmap, which speeds up the FIB initialization process. This method of using bitmaps is the same as the technique used for Neighbor Discovery and Prefix Delegation entries for IPv6 prefixes. Also, design enhancements in the queueing for packets directed to the CPU. Incoming ARP requests and packets to directly-connected destinations with unresolved ARP and packets to unknown destinations (that matches the catch-all entry and is forwarded to the CPU) are sent to queue Q0, instead being sent to the CPU on the same queue (Q5). Such a method of operation caused a delay in resolution of the address using ARP. Because the S6000 platform has more number of queues than the other S-Series platforms, such as S4810, packets triggering ARP resolution and packets destined or sent to the catch-all entry are forwarded to queue Q0, and ARP request packets are transmitted to queue Q5. This method of servicing packets also applies for IPv6 traffic.

RDMA Over Converged Ethernet (RoCE) Overview

This functionality is supported on the S6000 platform.

Remote direct memory access (RDMA) is a mechanism that reduces both CPU cycles and latency. RDMA over converged Ethernet (RoCE) implements IB over Ethernet. RRoCE sends InfiniBand (IB) packets over IP. IB supports input and output connectivity for the Internet infrastructure. InfiniBand is supported to enable the expansion of network topologies over large geographical boundaries and creation of next-generation I/O interconnect standard in servers. Although the endpoints or the destination servers generate such RRoCE packets, from the perspective of the switch, RRoCE is considered and processed as an IP packet.

RRoCE packets are received and transmitted on specific interfaces called lite-subinterfaces. These interfaces are similar the normal L3 physical interfaces with the exception of additional provisioning that they offer to enable the VLAN ID for encapsulation.

You can configure a physical interface or a L3 Port Channel interface as a lite-subinterface. When you configure a lite subinterface, only tagged IP packets with VLAN encapsulation are processed and routed. All other data packets are discarded.

To provide lossless service for RRoCE, Qos service policy must be configured in the ingress direction, such as dot1p and PFC and in the egress direction, such as strict priority for queues mapped to the VLAN dot1p values, on lite-subinterfaces.

Normal L3 physical interface processes only untagged packets and makes routing decisions based on the default L3 VLAN ID(4095), while the routed packets are transmitted as untagged.

To enable routing of RRoCE packets, the VLAN ID is mapped to the default VLAN ID of 4095 and this mapping is performed using VLAN translation. After VLAN translation, the RRoCE packets are considered in the same manner as normal IP packets that received on L3 interface and routed in the egress direction.

At the egress interface, the VLAN ID is appended to the packet and transmitted out of the interface as a tagged packet with the dot1Q value preserved.

RDMA is a technology using which a virtual machine (VM) can directly transfer information the memory of another VM, thereby enabling VMs to be connected to storage networks. With RoCE, RDMA enables data to be forwarded without passing through the CPU and the main memory path of TCP/IP. In a deployment that contains the RoCE network and the normal IP network, called backend and front-end network segments respectively, on two different networks, RRoCE enables the RoCE and the regular IP networks to be combined and RoCE frames to be sent over the IP network. This method of transmission, called RRoCE, results in the encapsulation of RoCE packets to IP packets.

When a storage area network (SAN) is connected over an IP network, the following conditions must be satisfied:

- Faster Connectivity: QoS for RRoCE enables faster and lossless nature of disk input and output services.
- Lossless connectivity: VMs require the connectivity to the storage network to be lossless at all times. When an upgrade of the network nodes is performed in a planned manner, especially with top-of-rack (ToR) nodes where there is a single point of failure for the VMs, disk I/O operations are expected to occur in 20 seconds. If disk in not accessible in 20 seconds, unexpected and undefined behavior of the VMs occurs. You can enable the optimization mechanism for booting time of the ToR nodes that experience a single point of failure to reduce the outage in traffic-handling operations to be less.

RoCE over a routed system is called RRoCE. RRoCE has IP headers. RRoCE is bursty and uses the entire 10 Gigabit Ethernet interface. Although RRoCE and normal data traffic are propagated in separate network portions, it might also be necessary in certain topologies to combine both the RRoCE and data traffic in a single network structure. RRoCE traffic is marked with dot1p priorities 3 and 4 (code points 011 and 100, respectively) and these queues are strict and lossless. DSCP code points are not tagged for RRoCE. Both ECN and PFC are enabled for RRoCE traffic. For normal IP or data traffic that is not RRoCE-enabled, the packets comprise TCP and UDP packets and they can be marked with DSCP code points. Multicast is not supported in that network.

Preserving 802.1Q VLAN Tag Value for Lite Subinterfaces

This functionality is supported on the S6000 platform.

All the frames in a Layer 2 VLAN are identified using a tag defined in the IEEE 802.1Q standard to determine the VLAN to which the frames or traffic are relevant or associated. Such frames are encapsulated with the 802.1Q tags. If a single VLAN is configured in a network topology, all the traffic packets contain the same do1q tag, which is the tag value of the 802.1Q header. If a VLAN is split into multiple, different sub-VLANs, each VLAN is denoted by a unique 8021.Q tag to enable the nodes that receive the traffic frames determine the VLAN for which the frames are destined.

Typically, a L3 physical interface processes only untagged or priority-tagged packets. The routing decision is made based on the default L3 VLAN ID (4095) and routed packets are transmitted as untagged packets. Tagged packets that are received on L3 physical interfaces are dropped. To enable the routing of tagged packets, the port that receives such tagged packets needs to be configured as a switchport and must be bound to a VLAN as a tagged member port.

A lite-subinterface is similar a normal L3 physical interface, except that additional provisioning is performed to set the VLAN ID for encapsulation. This setting is mainly used for data-plane routing of RRoCE packets.

A physical interface or a Layer 3 Port channel interface can be configured as a lite-subinterface. Once a lite-subinterface is configured, only tagged IP packets with encapsulation VLAN are processed and routed. All other data packets are discarded except the L2 and L3 control frames. It is not required for a

VLAN ID to be preserved (in the hardware or the OS application) when a VLAN ID used for encapsulation is associated with a physical/Port-channel interface. Normal VLANs and VLAN encapsulation can exist simultaneously and any non-unicast traffic received on a normal vlan is not flooded using lite-subinterfaces whose encapsulation VLAN ID matches with that of the normal VLAN ID.

You can use the `encapsulation dot1q vlan-id` command in INTERFACE mode to configure lite-sub-interfaces.

encapsulation dot1q

Configures lite-subinterfaces. This command is supported on the S6000 platform.

Syntax

`encapsulation dot1q vlan-id`

To remove a previously configured lite-subinterface, use the `no` version of this command.

Parameters

dot1q *vlan-id*	Enter the keyword `dot1q` followed by the VLAN ID to which the host belongs. The range is from 1 to 4094. A lite subinterface is considered as a Layer 3 port property and is synchronous with the existing rules of applying Layer 2 or Layer 3 properties to an interface.

Command Modes

INTERFACE

Command History

Version 9.3.0.0	Introduced on the S6000 platform.

Usage Information

To enable routing of RRoCE packets, the VLAN ID is mapped to the default VLAN ID of 4095 and this mapping is performed using VLAN translation. After VLAN translation, the RRoCE packets are considered in the same manner as normal IP packets that received on L3 interface and routed in the egress direction. At the egress interface, the VLAN ID is appended to the packet and transmitted out of the interface as a tagged packet with the dot1Q value preserved. The dot1Q value is preserved only for egress interfaces that are associated with a VLAN or a lite-subinterface . If a Layer 3 interface is configured without the encapsulation 802.1Q VLAN ID or is an untagged interface in a VLAN , the dot1Q value is not preserved .

9

Interfaces

This chapter describes the interfaces-related enhancements and contains the following sections:

- Enabling the Management Address TLV on All Interfaces of an Aggregator
- Enhanced Validation of Interface Ranges

Enabling the Management Address TLV on All Interfaces of an Aggregator

The management address TLV, which is an optional TLV and is of type 8, that denotes the network address of the management interface, is supported by the Dell Networking OS and it is advertised on all of the interfaces on an I/O Aggregator in the Link Layer Discovery Protocol (LLDP) data units. You can use the `show running-configuration` command to verify that this TLV is advertised on all of the configured interfaces and the `show lldp neighbors detail` command to view the value of this TLV.

Enhanced Validation of Interface Ranges

This functionality is supported on the S4810, S4820T, S6000, Z9000, MXL, and I/O Aggregator platforms.

On the S4810, S4820T, S6000, Z9000, MXL, and I/O Aggregator platforms, you can avoid specifying spaces between the range of interfaces separated by commas that you configure by using the `interface range` command. For example, if you enter a list of interface ranges, such as `interface range fo 2/0-1,te 10/0,gi 3/0,fa 0/0`, this configuration is considered valid. In releases of Dell Networking OS earlier than Release 9.3.0.0, if you enter such a combination of interfaces as a range in a comma-separated list, without spaces separating the ranges, an error message was displayed. Starting with Release 9.3.0.0, the comma-separated list is not required to be separated by spaces in between the ranges. You can associate multicast MAC or hardware addresses to an interface range and VLANs by using the `mac-address-table static` *`multicast-mac-address`* `vlan` *`vlan-id`* `output-range` *`interface`* command.

10

IPv4 Routing

This chapter describes the IPv4 routing-related enhancements and contains the following sections:

- IPv4 Path MTU Discovery Overview
- Configuring the Duration to Establish a TCP Connection
- Using Loopback Address in ICMP Unreachable Messages

IPv4 Path MTU Discovery Overview

In common network topologies, hosts send large large volumes of data to other neighboring devices using IP packets. For effective utilization of network resources, enhanced performance, and easy reassembly of packets that are transmitted, devices attempt to forward packets from the origin to the endpoint of the network without the need of fragmentation as much as possible. The size of the packet that can be sent across each hop in the network path without being fragmented is called the path maximum transmission unit (PMTU). This value might vary for the same route between two devices, mainly over a public network, depending on the network load and speed, and it is not a symmetric, consistent value. This MTU size can also be different for various types of traffic sent from one host to the same endpoint.

Path MTU discovery (PMTD) is a mechanism that identifies the path MTU value between the sender and the receiver, and uses the determined value to transmit packets across the network. PMTD, as described in RFC 1191, denotes that the default byte size of an IP packet is 576. The IP and TCP portions of the frame constitute 40 bytes and the remaining 536 bytes form the data paylod. This packet size is called the maximum transmission unit (MTU) for IP4 frames. PMTD operates by containing the do not fragment (DF) bit set in the IP headers of outgoing packets. When any device along the network path contains an MTU that is smaller than the size of the packet that it receives, the device drops the packet and sends an Internet Control Message Protocol (ICMP) Fragmentation Needed (Type 3, Code 4) message with its MTU value contained in the message to the source or the sending device. This message enables the source to identify that the transmitted packet size must be reduced. The packet is retransmitted with a lower size than the previous value. This process is repeated in an interative way until the MTU of the transmitted packet is lower or equal to the MTU of the receiving device for it to obtain the packet without fragmentation. If the ICMP message from the receiving device that is sent to the originating device contains the next-hop MTU, then the sending device lowers the packet size accordingly and resends the packet. Otherwise, the iterative method is followed until the packet can traverse without being fragmented.

PMTD is enabled by default on the switches that support this capability. To enable PMTD to function correctly, you must enter the `ip unreachables` command on a VLAN interface to enable the generation of ICMP unreachable messages. PMTD is supported on all the layer 3 VLAN interfaces. Because all of the Layer 3 interfaces are mapped to the VLAN ID of 4095 when VLAN subinterfaces are configured on it, it is not possible to configure unique layer 3 MTU values for each of the layer 3 interfaces. If a VLAN interface contains both IPv4 and IPv6 addresses configured on it, both the IPv4 and IPv6 traffic are applied the same MTU size; you cannot specify different MTU values for IPv4 and IPv6 packets. This functionality is supported on the S4810, S4820T, Z9000, and MXL platforms.

Using the Configured Source IP Address in ICMP Messages

This functionality to enable ICMP messages, such as ICMP unreachable or ICMP error messages, to be sent with the configured ICMP source interface IP address address instead of the front-end port IP address for traceroute command output listing is supported on the S4810, S4820T, Z9000, S6000, and MXL platforms.

ICMP error or unreachable messages are now sent with the configured IP address of the source interface, such as the loopback address of the system, instead of the front-end port IP address as the source IP address. This behavior is applicable if you enable the generation of ICMP unreachable messages by entering the `ip unreachable` command in Interface Configuration mode. When a ping or traceroute packet from an endpoint or a device arrives at the null 0 interface configured with a static route, it is discarded. In such cases, you can configure Internet Control Message Protocol (ICMP) unreachable messages to be sent to the transmitting device or the origin for such discarded packets. The null interface is a data sink that handles undesired traffic sent to a device because it does not forward or receive packets, and merely discards them and sends ICMP unreachable messages.

Configuring the ICMP Source Interface

In network environments that contain a large number of devices, ranging up to thousands of systems, and with each device configured for equal-cost multipath (ECMP) links, you cannot effectively and optimally use the traceroute and ping applications to examine the network reachability and identify any broken links for diagnostic purposes. In such cases, if the reply that is obtained from each hop on the network path contains the IP address of the adjacent, neighboring interface from which the packet is received, it is difficult to employ the ping and traceroute utilites. You can enable the ICMP error and unreachable messages to contain the configured IP address of the source device instead of the previous hop's IP address to be able to easily and quickly identify the device and devices along the path because the DNS server maps the loopback IP address to the hostname and does not translate the IP address of every interface of the switch to the hostname.

You can enable the mechanism to configure the source or the originating interface from which the packet (the device that generates the ICMP error messages) is received by the switch to send the configured source interface IP address instead of its front-end IP address to be used in the ICMP unreachable messages and in the `traceroute` command output. You can use the `ip icmp source-interface` *`interface`* or the `ipv6 icmp source-interface` *`interface`* commands in Configuration mode for IPv4 and IPv6 packets respectively to enable the ICMP error messages to be sent with the source interface IP address. This functionality is supported on loopback, VLAN, port channel, and physical interfaces for IPv4 and IPv6 messages. This capability to configure the source interface to send the IP address is not supported on tunnel interfaces. ICMP error relay, PATH MTU transmission, and fragmented packets are not supported for tunnel interfaces. The traceroute utilities, for IPv4 and IPv6, list the IP addresses of the devices in the hops of the path for which ICMP source-interface is configured.

Working of the Traceroute Utility

Traceroute sends a sequence of three ICMP echo request packets addressed to a destination host. The time-to-live (TTL) value, also known as hop limit, is used in determining the intermediate routers being traversed towards the destination. Routers decrement packets' TTL value by 1 when routing and discard packets whose TTL value has reached zero, returning the ICMP error message, ICMP Time Exceeded. Common default values for TTL are 128 (Windows OS) and 64 (Unix-based OS).

Traceroute works by sending packets with gradually increasing TTL value, starting with TTL value = 1. The first router receives the packet, decrements the TTL value and drops the packet because it then has a TTL value of zero. The router sends an ICMP Time Exceeded message back to the source. The next set of packets are given a TTL value of 2; therefore, the first router forwards the packets, but the second router drops them and replies with ICMP Time Exceeded. With such a progressive pattern, the traceroute application uses the returned ICMP Time Exceeded messages to build a list of routers that packets traverse, until the destination is reached and returns an ICMP Echo Reply message.

On Unix-based operating systems, the traceroute utility uses User Datagram Protocol (UDP) datagrams by default, with destination port numbers ranging from 33434 to 33534.

ip icmp source-interface

Enable the ICMP error and unreachable messages to be sent with the source interface IP address, such as the loopback address, instead of the hops of the preceding devices along the network path to be used for easy debugging and diagnosis of network disconnections and reachability problems with IPv4 packets. This functionality is supported on the S4810, S4820T, S6000, Z9000, and MXL platforms.

Syntax `ip icmp source-interface` *interface*

Parameters

interface Enter one of the following keywords and slot/port or number information:

- For a Management Ethernet interface, enter the keyword `managementethernet`.

  NOTE: When you configure the capability to enable the loopback IP address to be sent for easy debugging and diagnosis (IP addresses of the devices for which the ICMP source interface is configured), the source IP address of the outgoing ICMP error message is modified, although the packets are not sent out using the configured interface. Because the management interface is configured without any parameters such as the IP address, it is treated to the management interface of the primary unit or the existing unit.
- For a Loopback interface, enter the keyword `loopback`. The range is from 0 to 16383.
- For a 1-Gigabit Ethernet interface, enter the keyword `GigabitEthernet`.
- For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet`.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE`.
- For a VLAN interface, enter the keyword `vlan`. The range is from 1 to 4094.

Defaults Not configured.

Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, S6000, Z9000, and MXL platforms.

Usage Information

You can enable the mechanism to configure the source or the originating interface from which the packet (the device that generates the ICMP error messages) is received by the switch to send the loopback address instead of its source IP address to be used in the ICMP unreachable messages and in the `traceroute` command output. The loopback address must be unique in a particular domain.

In network environments that contain a large number of devices, ranging up to thousands of systems, and with each device configured for equal-cost multipath (ECMP) links, you cannot effectively and optimally use the traceroute and ping applications to examine the network reachablity and identify any broken links for diagnostic purposes. In such cases, if the reply that is obtained from each hop on the network path contains the IP address of the adjacent, neighboring interface from which the packet is received, it is difficult to employ the ping and traceroute utilites. You can enable the ICMP unreachable messages to contain the loopback address of the source device instead of the previous hop's IP address to be able to easily and quickly identify the device and devices along the path because the DNS server maps the loopback IP address to the hostname and does not translate the IP address of every interface of the switch to the hostname.

Example

```
FTOS(conf)#ip icmp source-interface tengigabitethernet 0/0
FTOS(conf)#
```

ipv6 icmp source-interface

Enable the ICMP error and unreachable messages to be sent with the source interface IP address, such as the loopback address, instead of the hops of the preceding devices along the network path to be used for easy debugging and diagnosis of network disconnections and reachability problems with IPv6 packets. This functionality is supported on the S4810, S4820T, S6000, Z9000, and MXL platforms.

Syntax `ipv6 icmp source-interface` *interface*

Parameters

interface	Enter one of the following keywords and slot/port or number information: • For a Management Ethernet interface, enter the keyword `managementethernet`.

NOTE: When you configure the capability to enable the loopback IP address to be sent for easy debugging and diagnosis (IP addresses of the devices for which the ICMP source interface is configured), the source IP address of the outgoing ICMP error message is modified, although the packets are not sent out using the configured interface. Because the management interface is configurable only without any parameters such as the IP address, it is treated to the management interface of the primary unit or the existing unit.

- For a Loopback interface, enter the keyword `loopback`. The range is from 0 to 16383.
- For a 1-Gigabit Ethernet interface, enter the keyword `GigabitEthernet`.
- For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet`.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE`.
- For a VLAN interface, enter the keyword `vlan`. The range is from 1 to 4094.

Defaults Not configured.

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the S4810, S4820T, S6000, Z9000, and MXL platforms.

Usage Information

You can enable the mechanism to configure the source or the originating interface from which the packet (the device that generates the ICMP error messages) is received by the switch to send the loopback address instead of its source IP address to be used in the ICMP unreachable messages and in the `traceroute` command output. The loopback address must be unique in a particular domain.

In network environments that contain a large number of devices, ranging up to thousands of systems, and with each device configured for equal-cost multipath (ECMP) links, you cannot effectively and optimally use the traceroute and ping applications to examine the network reachablity and identify any broken links for diagnostic purposes. In such cases, if the reply that is obtained from each hop on the network path contains the IP address of the adjacent, neighboring interface from which the packet is received, it is difficult to employ the ping and traceroute utilites. You can enable the ICMP unreachable messages to contain the loopback address of the source device instead of the previous hop's IP address to be able to easily and quickly identify the device and devices along the path because the DNS server maps the loopback IP address to the hostname and does not translate the IP address of every interface of the switch to the hostname.

Example	`FTOS(conf)#ipv6 icmp source-interface tengigabitethernet 0/0` `FTOS(conf)#`

Configuring the Duration to Establish a TCP Connection

This procedure is supported on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

You can configure the amount of time for which the device must wait before it attempts to establish a TCP connection. Using this capability, you can limit the wait times for TCP connection requests. Upon responding to the initial SYN packet that requests a connection to the router for a specific service (such as SSH or BGP) with a SYN ACK, the router waits for a period of time for the ACK packet to be sent from the requesting host that will establish the TCP connection.

You can set this duration or interval for which the TCP connection waits to be established to a significantly high value to prevent the device from an out-of-service condition or becoming unresponsive during a SYN flood attack that occurs on the device. You can set the wait time to be 10 seconds or lower. If the device does not contain any BGP connections with BGP neighbors across WAN links, you must set this interval to a higher, appropriate value, depending on the complexity of your network and the configuration attributes.

To configure the duration for which the device waits for the ACK packet to be sent from the requesting host to establish the TCP connection, perform the following steps:

1. Define the wait duration in seconds for the TCP connection to be established.

   CONFIGURATION mode

   `Dell(conf)#ip tcp reduced-syn-ack-wait <9-75>`

   You can use the `no ip tcp reduced-syn-ack-wait` command to restore the default behavior, which causes the wait period to be set as 8 seconds.
2. View the interval that you configured for the device to wait before the TCP connection is attempted to be established.

   EXEC mode

   `Dell>show ip tcp reduced-syn-ack-wait`

ip tcp initial-time

Define the wait duration in seconds for the TCP connection to be established. This command is supported on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

Syntax	`ip tcp initial-time <8-75>` To restore the default behavior, which causes the wait period to be set as 8 seconds, use the `no ip tcp initial-time` command	
Parameters	***<8-75>***	Wait duration in seconds for the TCP connection to be established.
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

Usage Information	You can configure the amount of time for which the device must wait before it attempts to establish a TCP connection. Using this capability, you can limit the wait times for TCP connection requests. Upon responding to the initial SYN packet that requests a connection to the router for a specific service (such as SSH or BGP) with a SYN ACK, the router waits for a period of time for the ACK packet to be sent from the requesting host that will establish the TCP connection.

show ip tcp initial-time

Displays the interval that you configured for the device to wait before the TCP connection is attempted to be established. This command is supported on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

Syntax	`show ip tcp initial-time`	
Command Modes	EXEC EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

11

Link Aggregation Groups (LAGs)

This chapter describes the link aggregation control protocol (LACP) and link aggregation group (LAG) enhancements and contains the following sections:

- Configuring the Uplink Speed of Interfaces as 40 Gigabit Ethernet
- Configuring the Minimum Number of Links to be Up for Uplink LAGs to be Active
- Optimizing Traffic Disruption Over LAG Interfaces On IOA Switches in VLT Mode
- Preserving LAG and Port Channel Settings in Nonvolatile Storage
- Setting Up a Threshold for Utilization of High-Gigabit Port Channels
- Monitoring the Member Links of a LAG Bundle
- Enabling the Verification of Member Links Utilization in a High-Gigabit Port Channel

Configuring the Minimum Number of Links to be Up for Uplink LAGs to be Active

You can enable the mechanism to activate the LAG bundle for uplink interfaces or ports (the uplink port-channel (LAG 128)) on the I/O Aggregator only when a minimum or the required number of member interfaces of the LAG bundle are up. For example, based on your network deployment, you you might want the uplink LAG bundle to be activated only if a certain number of member interface links are also in the administratively up state. If you enable this setting, the uplink LAG bundle is brought up only when a specified minimum number of links are up and the LAG bundle is moved to the down state when the number of active links in the LAG becomes less than the specified minimum number of interfaces. By default, the uplink LAG 128 interface is activated when at least one member interface is up. The Aggregator brings down the LAG bundle when the number of operational links drops below the specified minimum number and automatically brings up the LAG bundle when the number of operational links equals or exceeds the configured minimum number.

To configure the minimum or the least number of member links of a LAG bundle or a port channel that must be up for a LAG bundle to be fully up, perform the following:

Specify the minimum number of member interfaces of the uplink LAG 128 bundle that must be up for the LAG bundle to be brought up. The default minimum number of member links that must be active for the uplink LAG to be active is 1. You can enter the `minimum-links number` command in the Port Channel Interface 128 Configuration mode to specify this value.

PORT-CHANNEL INTERFACE 128 (conf-if-po-128)

```
Dell(conf-if-po-128)#minimum-links 4
```

You can use the `show interfaces port-channel` command to view information regarding the configured LAG or port channel settings. The **Minimum number of links to bring Port-channel up is** field in the output of this command displays the configured minimum number of active links for the LAG to be enabled.

Optimizing Traffic Disruption Over LAG Interfaces On IOA Switches in VLT Mode

When an Aggregator operates in VLT mode, the VLT LAG configurations are saved in nonvolatile storage (NVS) when you enter the write memory command. This method of saving the VLT settings in NVS and restoring the saved settings when the Aggregator restarts reduces the disruption of traffic that is handled during the restart of the primary and secondary VLT peer nodes. By restoring the settings saved in NVS, the VLT ports come up in a quicker way on the primary VLT node. On the secondary VLT peer node, the delay in restoration of the VLT LAG parameters is reduced (90 seconds by default) before it becomes operationally up. This makes sure that the configuration settings of the primary VLT node are synchronized with the secondary VLT peer node before the secondary VLT mode becomes operational.

In VLT mode of the Aggregator, the LAG bundle is automatically established when the LACP PDU is received on the port and LAG configurations are not saved. When the VLT LAG parameters are not saved in NVS, this behavior impacts the VLT port from being brought up in a faster, effective manner during the VLT node restart because of the switchover of traffic from the primary to the secondary VLT peer node occurring before the restarted peer node is synchronized with the other peer node.

The VLT domain details, such as the domain ID, the port-channel number that functions as the VLT interconnect link, the default MAC address of the domain, and the unique unit ID of each peer in the VLT domain, are stored in NVS. Also, the port-channel member interface details and the VLT LAG configuration are stored. The traffic outage is less than 200 millisconds during the restart or switchover of the VLT peer nodes from primary to secondary.

Preserving LAG and Port Channel Settings in Nonvolatile Storage

You can now use the `write memory` command on an I/O Aggregator that operates in standalone and stacking modes, which saves the running configuration to the startup configuration to be preserved across reboots of the device, to save the LAG port channel configuration parameters. All of the statically configured port channels (LAG 128) and automatically-configured internal LAGs are saved. This behavior enables the port channels to be brought up in a faster way because the interface attributes that are configured are available in the system database during the booting of the device. With the reduction in time for the port channels to become active after the switch is booted, the loss in the number of packets that are serviced by these interfaces is minimized.

Enabling the Verification of Member Links Utilization in a LAG Bundle

To examine the working efficiency of the LAG bundle interfaces, do the following

1. The functionality to detect the working efficiency of the LAG bundle interfaces is automatically activated on all the port channels, except the port channel that is configured as a VLT interconnect link, during the booting of the switch.
2. Use the `show link-bundle-distribution port-channel` *`interface-number`* to display the traffic-handling and utilization of the member interfaces of the port channel. The following sample output is displayed when you enter this `show` command.

   EXEC

   ```
   Dell#show link-bundle-distribution port-channel
   ```

   ```
   Dell#show link-bundle-distribution port-channel 1
   Link-bundle trigger threshold - 60

   LAG bundle - 1      Utilization[In Percent] - 0      Alarm State - Inactive

   Interface                Line Protocol   Utilization[In Percent]
   Te 0/5                   Up              0
   Te 0/13                  Up              0
   ```

Monitoring the Member Links of a LAG Bundle

You can examine and view the operating efficiency and the traffic-handling capacity of member interfaces of a LAG or port channel bundle. Such a method of analyzing and tracking the number of packets processed by the member interfaces in percentage enables you to optimally and effectively manage and distribute the packets that are handled by the LAG bundle. The functionality to detect the working efficiency of the LAG bundle interfaces is automatically activated on all the port channels, except the port channel that is configured as a VLT interconnect link, during the booting of the switch. This mechanism is supported on I/O Aggregators in stacking, standalone, and VLT modes and it is not supported in programmable MUX (PMUX) mode. By default, this capability is enabled on all of the port channels set up on the switch.

You can use the `show link-bundle-distribution port-channel` *`interface-number`* to display the traffic-handling and utilization of the member interfaces of the port channel. The following table describes the output fields of this show command.

Table 3. Output Field Descriptions for show link-bundle-distribution port-channel Command

Field	Description
Link-bundle trigger threshold	Threshold value that is the checkpoint, exceeding which the link bundle is marked as being overutilized and alarm is generated
LAG bundle number	Number of the LAG bundle
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the port channel
Alarm State	Indicates whether an alarm is generated if overutilization of the port channel occurred. The value, Active, is displayed for this field.

Field	Description
Interface	Slot and port number, and the type of the member interface of the port channel
Line Protocol	Indicates whether the interface is administratively up or down
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the particular member interface

You can also use the `show running-configuration interface port-channel` command in EXEC Privilege mode to view whether the mechanism to evaluate the utilization of the member interfaces of the LAG bundle is enabled. The following sample output illustrates the portion of this `show` command:

```
Dell#show running-config int port-channel
!
interface Port-channel 1
 mtu 12000
 portmode hybrid
 switchport
 vlt-peer-lag port-channel 1
 no shutdown
 link-bundle-monitor enable
```

show link-bundle-distribution port-channel

Display the traffic-handling and utilization of the member interfaces of the port channel.

Syntax `show link-bundle-distribution port-channel` *interface-number*

Parameters

interface-number	For a Port Channel interface, enter the keyword `port-channel` followed by a number: Range: 1-128

Command Modes

EXEC

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator

Usage Information

The following table describes the output fields of this `show` command:

Field	Description
Link-bundle trigger threshold	Threshold value that is the checkpoint, exceeding which the link bundle is marked as being overutilized and alarm is generated
LAG bundle number	Number of the LAG bundle
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the port channel

Field	Description
Alarm State	Indicates whether an alarm is generated if overutilization of the port channel occurred. Possible values are Active and Inactive
Interface	Slot and port number, and the type of the member interface of the port channel
Line Protocol	Indicates whether the interface is administratively up or down
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the particular member interface

Example

```
Dell#show link-bundle-distribution port-channel 1
Link-bundle trigger threshold - 60

LAG bundle - 1      Utilization[In Percent] - 0      Alarm
State - Inactive

Interface               Line Protocol   Utilization[In
Percent]
Te 0/5                  Up              0
Te 0/13                 Up              0
```

Setting Up a Threshold for Utilization of High-Gigabit Port Channels

This functionality is supported on the Z9000 platform.

You can configure a mechanism to monitor a backplane high-Gigabit Ethernet port channel and generate a system logging message or an SNMP trap when the traffic distribution and the handled data packets on the bundle is uneven or inconsistent. The formula or the computation parameter to determine the uneven or the unequal distribution of traffic is predefined and at a particular point in time, if you enable the capability to examine the efficiency of the member links of a port channel bundle, such an unbalanced segregation of traffic across the member links of the high-Gigabit Ethernet bundle is indicated using alarms and traps. Also, when the traffic is resumed to be handled in an equalized, proper manner, a notification using alarms and SNMP traps is generated.

The Dell Networking OS already contains the functionality to monitor the performance and traffic-handling of virtual interfaces created as LAG bundles and ECMP configured on physical user ports. You can now verify the traffic-distribution and processing of high-Gigabit Ethernet port channels. Trunk groups for backplane higig link bundles between leaf and spines are created. For trunk groups to be provisioned on Z9000 platforms, 1 trunk group (hiGig link bundle) on each leaf unit is created and 4 trunk groups on each spine unit are created. As a result, a total of 12 trunk groups are present on the 2 spine and 4 leafs of Z9000 platform.

Based on the hashing algorithm that is specified, traffic in the trunk groups are distributed. It is possible that an unequal or imbalanced traffic distribution in higig trunk groups might occur. When you configure the method to monitor the high-Gigabit Ethernet port channel and trunk groups, you can view and analyze the unequal traffic split and flow in the trunk groups and take corrective action as appropriate.

You can use this optimal, cohesive capability in your network environment to detect whether the configured applications or utilities are causing traffic to be unevenly distributed on a higig link bundle for best performance. This capability to monitor the port channel bundles is applicable for any platform that contains backplane high-Gigabit Ethernet links.

The collected and derived data rates for the configured rate-interval monitor and examine the working efficiency and traffic-handling capacity of the LAG bundles on high-Gigabit Ethernet trunk interfaces that are created statically.

You can use the mechanism to examine the working efficiency of the LAG bundle interfaces to adjust and modify the switch for effective utilization of backplane links.

Alarms are generated if the port-channel threshold is greater than the configured threshold and the unevenness is greater than 10 percent between links for three successive rate-intervals.

Alarms are removed if the port-channel threshold becomes lower than the configured threshold and the unevenness is less than 10 percent between links for three successive rate-interval, multiplied by 3 time intervals.

The following log messages are generated when the threshold for high-Gigabit port channel or LAG bundle monitoring is exceeded:

- An informational message when an alarm is triggered for uneven distribution observed in a LAG bundle
- An informational message when the alarm is cleared

The following additional information is recorded in the alarm, apart from the usual, standard details such as module name and timestamp of generation:

- Link bundle name (hg-port-channel slot/NpuId/BundleId)
- Alarm raising or clearing

The following examples display the system log messages triggered when the threshold for high-Gigabit port channel monitoring is exceeded:

- **%STKUNIT0-M:CP %SWMGR-5-HG-BUNDLE_UNEVEN_DISTRIBUTION: Found uneven distribution in hg-port-channel 0/5/0**
- **%STKUNIT0-M:CP %SWMGR-5-HG-BUNDLE_UNEVEN_DISTRIBUTION_ALARM_CLEAR: Uneven distribution in hg-port-channel 0/5/0 got cleared**

Guidelines for Configuring the Mechanism to Monitor High-Gigabit Port Channels

Keep the following points in mind when you activate and configure the capability to examine the utilization and working-efficiency of backplane high-Gigabit Ethernet port channels as trunk groups:

- By default, the capability to monitor the traffic utilization and distribution of high-Gigabit Ethernet trunk groups is disabled.
- Because each NPU unit in each line card (or control processor card) can contain multiple trunk groups (high-Gigabit port channels). The interface specifier convention for hg-port-channel is slotId/npuUnitId/localPortChannelId, which denote the slot, NPU, and the port channel identifiers.
- For Z9000, slotId (stack unitId) is constant and does not vary. NpuUnitId ranges from 0-5 and local portChannelId ranges from 0-0 for leaf NpuUnits and 0-3 for spine NpuUnits.
- Link-bundle monitoring is commenced if monitoring is enabled for the bundle and when the bundle egress utilization exceeds a threshold. This behavior is required if you want to view the utilization

alarms only when the utilization levels are high. At low utilization levels, it is possible that there are only one or two significant flows can cause unevenness. Such an imbalanced traffic flow is not critical or indicative of a problem. The higig link bundle trigger threshold is a system-wide or a global setting for the device.

- If you enabled the generation of SNMP traps, syslogs and traps are transmitted when an uneven distribution is observed. Another syslog and trap is generated when the unevenness is cleared.
- Link bundle utilization is calculated as the bandwidth-weighted mean utilization of all links in a bundle (calculated as [total bandwidth of all links / total bytes-per-sec of all links]). This calculation is performed only on those links that are up on their operational status.
- Rate Interval for polling the traffic statistics for member links of the high-Gigabit port channel needs to be configured. The default hiGig stats polling interval is 15 seconds. This interval cannot be configured per high-Gigabit port channel and is applicable for all of the high-Gigabit port channels on the system.
- The treshold value identifies when to start the link bundle utilization calculation trigger (default of 60 percent). When overall utilization (mean) is below this value, link bundle distribution unevenness will not be reported.
- If unevenness is observed over 3 consecutive measurements, an alarm event shall be generated. The time interval between 2 measurements is defined by the rateInterval for Hg stats polling (default 15 seconds). Alarm clear is sent when evenness is observed for three successive rate intervals. If individual link utilization information is not available for a given timestamp, link bundle utilization will not be calculated at that time stamp. The previous known record shall be used for the alarm calculation.
- Turning on and off the link bundle monitoring is performed at a high-Gigabit Ethernet port-channel level configuration.
- The difference of utilization % between the high used link and low used link is used to determine the alarm condition. The alarm trigger reporting is based on the same algorithm used for link bundle monitoring on LAG/ECMP. The alarm reporting is triggered when the configured threshold is crossed for a given bundle. At this time, if the delta utilization is beyond 10%, alarm is raised. The alarm condition remains active until all interface utilizations are within the band or until the overall utilization goes below the trigger threshold. An alarm is not raised or cleared instantaneously

Enabling the Verification of Member Links Utilization in a High-Gigabit Port Channel

This procedure is supported on the Z9000 platform.

To activate the mechanism to examine the working efficiency of the high-Gigabit Ethernet port channel interfaces, do the following

1. Use the `hg-link-bundle-monitor slot` *`slotId`* `npuUnit` *`npuUnitId`* `hg-port-channel` *`portChannelId`* `enable` command in Global Configuration mode to enable this functionality to detect the working efficiency of the high-Gigabit port channel bundle interfaces.

   CONFIGURATION mode

   ```
   Dell(conf)#hg-link-bundle-monitor slot 0 npuUnit 0 hg-port-channel 0 enable
   ```

2. Specify the trigger threshold for higig link bundle monitoring.

   CONFIGURATION mode

   ```
   Dell(conf)#hg-link-bundle-monitor trigger-threshold 30
   ```

3. Specify the interval in seconds for higig link bundle monitoring.
   CONFIGURATION mode
   ```
   Dell(conf)#hg-link-bundle-monitor rate-interval 10
   ```
4. Enable the generation of traps for higig link-bundle monitoring.
   CONFIGURATION mode
   ```
   Dell(conf)#snmp-server enable traps hg-lbm
   ```
5. Use the `show hg-link-bundle-distribution` command to display the traffic-handling and utilization of the member interfaces of the port channel. The following table describes the output fields of this show command.
   EXEC, EXEC Privilege modes
   ```
   Dell#show hg-link-bundle-distribution 0 npuUnit 5 hg-port-channel 0
   ```

hg-link-bundle-monitor

Enable the capability to examine the utilization and traffic distribution of high-Gigabit port channels. This command is supported on the Z9000 platform.

Syntax

`hg-link-bundle-monitor slot` *`slotId`* `npuUnit` *`npuUnitId`* `hg-port-channel` *`portChannelId`* `enable`

To disable this capability, use the `no` version of this command..

Parameters

slot *slotId*	Enter the keyword `slot` followed by the slot ID of the high-Gigabit port channel. For Z9000, the only valid slot number is 0.
npuUnit *npuUnitId*	Enter the keyword `npuUnit` followed by the NPU value. The range is from 0-5.
hg-port-channel *portChannelId*	Enter the keyword `hg-port-channel` followed by the unique ID of the port channel. Number of hg-port-channels vary for switch NPU and fabric NPUs.
enable	Enable the capability to examine the utilization and traffic distribution of high-Gigabit port channels.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the Z9000 platform.

Usage Information

You can configure a mechanism to monitor a backplane high-Gigabit Ethernet port channel bundle gig link bundle and generate a system logging message or an SNMP trap when the traffic distribution and the handled data packets on the bundle is uneven or inconsistent. The formula or the computation parameter to determine the uneven or the inequal distribution of traffic is predefined and at a particular point in time, if you enable the capability to examine the efficiency of the member links of a port channel bundle, such an unbalanced segregation of traffic across the member links of the high-Gigabit Ethernet bundle is indicated using alarms and

traps. Also, when the traffic is resumed to be handled in an equalized, proper manner, a notification using alarms and SNMP traps is generated.

hg-link-bundle-monitor trigger-threshold

Specify the threshold value for high-Gigabit Ethernet port channels or trunk groups, which is a checkpoint exceeding which the link bundle is marked as being overutilized and alarm is generated. This command is supported on the Z9000 platform.

Syntax

```
hg-link-bundle-monitor trigger-threshold <1-90>
```

To restore the default value, use the `no` version of this command..

Parameters

<1-90>	Trigger-threshold value in percentage

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the Z9000 platform.

Defaults

The default threshold value is 60.

Usage Information

Threshold for identifying when to start the link bundle utilization calculation trigger is fixed at a default of 60 percent. When overall utilization (mean) is below this value, link bundle distribution unevenness will not be reported. If unevenness is observed over 3 consecutive measurements, an alarm event shall be generated. The time interval between 2 measurements is defined by the rate interval for high statistics polling (default 15 seconds). Alarm clear is sent once evenness is observed for three successive rater interval periods. If individual link utilization information is not available for a given timestamp, link bundle utilization will not be calculated at that time stamp. The previous known record shall be used for the alarm calculation.

hg-link-bundle-monitor rate-interval

Specify the interval or frequency in seconds for polling the traffic statistics for member links of the high-Gigabit port channel. This command is supported on the Z9000 platform.

Syntax

```
hg-link-bundle-monitor rate-interval <10-299>
```

To restore the default value, use the `no` version of this command.

Parameters

<10-299>	Interface rate interval in seconds

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the Z9000 platform.

Defaults The default hiGig stats polling interval is 15 seconds.

Usage Information This interval cannot be configured per high-Gigabit port channel and is applicable for all of the high-Gigabit port channels on the system.

show hg-link-bundle-distribution

Display the traffic-handling and utilization of the member interfaces of the high-Gigabit port channel or trunk group. This command is supported on the Z9000 platform.

Syntax `show hg-link-bundle-distribution slot slotId npuUnit npuUnitId hg-port-channel portChannelId`

Parameters

slot *slotId*	Enter the keyword `slot` followed by the slot ID of the high-Gigabit port channel. For Z9000, the only valid slot number is 0.
npuUnit *npuUnitId*	Enter the keyword `npuUnit` followed by the NPU value. The range is from 0-5.
hg-port-channel *portChannelId*	Enter the keyword `hg-port-channel` followed by the unique ID of the port channel. Number of hg-port-channels vary for switch NPU and fabric NPUs.

Command Modes EXEC, EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the Z9000 platform.

Usage Information The following table illustrates the fields displayed in the output of this command:

Field	Description
Link-bundle trigger threshold	Threshold value that is the checkpoint, exceeding which the link bundle is marked as being overutilized and alarm is generated
Slot	Slot number where the high-Gigabit port-channel resides
npuUnit	Network Processign Unit (NPU) number where the high-Gigabit port-channel resides
number	Number of the LAG bundle
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the port channel
Alarm State	Indicates whether an alarm is generated if uneven utilization of the

Field	Description
	port channel occurred. Possible values are Active and Inactive
Interface	Slot and port number, and the type of the member interface of the port channel
Utilization (In Percent)	Traffic usage in percentage of the packets processed by the particular member interface

Example

```
FTOS#show hg-link-bundle-distribution 0 npuUnit 5 hg-port-
channel 0

hg-link-bundle trigger threshold - 60
Slot 0 npuUnit 5 hg-port-channel-0 Utilization [In Percent] -
0 Alarm State - Inactive
Interface Utilization [In Percent]
0/5:hg0           10
0/5:hg1           10
0/5:hg2           10
0/5:hg3           10
```

snmp-server enable traps (for High-Gigabit Port Channel)

Enable the generation of SNMP traps and notifications when the capability to examine the traffic utilization and distribution of high-Gigabit port channel links or trunk groups is enabled. This command is supported on the Z9000 platform.

Syntax

```
snmp-server enable traps [notification-type]
```

Parameters

notification type	Enter the keyword `hg-lbm` to enable high-Gigabit Link Bundle Monitoring traps

Command Modes

CONFIGURATION mode

Command History

Version 9.3.0.0	Introduced on the Z9000 platform.

show hardware stack-unit (for high-Gigabit Ethernet ports)

Display the data plane or management plane input and output statistics of the high-Gigabit Ethernet or backplane port of the designated stack unit or Z9000 unit.

NOTE: Only the parameters that are newly introduced with this command in Release 9.3(0.0) are explained here. For a complete description of all of the options that are available with this command, refer the relevant *Command Reference Guide* of the applicable platform of Release 9.2(0.0).

Z9000

Syntax

```
show hardware stack-unit stack-unit {buffer unit {0–5} [port
port-number]| cpu data-plane statistics | cpu i2c statistics |
cpu party-bus statistics | cpu sata-interface statistics |
drops [unit number [port port-number]] | hg-stats [unit number
[port port-number]] | ipmc-replication | stack-port port-number
| table-dump| unit unit-number {counters | details | port-stats
[detail] | register}}
```

Parameters

hg-stats [unit *unit-number* [port *port-number* | no]]	Enter the keyword `hg-stats` to display high–Gigabit Ethernet or backplane port buffer and queue statistics on the selected stack member. Optionally, use the keyword `unit` with a number to select port-pipe 0 to 5, and then use `port port-number` to select a port on that port-pipe. For Z9000, valid backplane ports for leaf NPU units (units 0–3) range from 34-41 and for spine NPU units (units 4–5) range from 1-16.

Defaults

none

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 9.3.0.0	Added support for the `hg-stats` option on the Z9000 platform.

Example (High-Gigabit Ethernet Statistics)

```
FTOS# show hardware stack-unit 0 hg-stats unit 4 port 30
% Error : Port 30 is not a valid back-plane hiGig port

FTOS# show hardware stack-unit 0 hg-stats unit 4 port 1
Input Statistics:
     3942277 packets, 4224329282 bytes
     0 64-byte pkts, 75905 over 64-byte pkts, 807091 over 127-
byte pkts
     300653 over 255-byte pkts, 245844 over 511-byte pkts,
2512784 over 1023-byte pkts
     394612 Multicasts, 0 Broadcasts
     0 runts, 0 giants, 0 throttles
     0 CRC, 0 overrun, 0 discarded
Output Statistics:
     1335309 packets, 187184751 bytes, 0 underruns
     0 64-byte pkts, 665971 over 64-byte pkts, 586727 over 127-
byte pkts
     82038 over 255-byte pkts, 58 over 511-byte pkts, 515 over
1023-byte pkts
     408949 Multicasts, 0 Broadcasts, 926163 Unicasts
     0 throttles, 0 discarded, 0 collisions, 0 wredDrops
Rate info (interval 30 seconds):
     Input 00.08 Mbits/sec,         10 packets/sec, 0.00% of
line-rate
     Output 00.00 Mbits/sec,          3 packets/sec, 0.00% of
line-rate
```

Related Commands	clear hardware system-flow — clears the statistics from selected hardware components. show interfaces stack-unit — displays information on all interfaces on a specific S-Series stack member. show processes cpu (S-Series) — displays the CPU usage information based on the processes running in an S-Series. show system (S-Series and Z-Series) — displays the current status of all the stack members or a specific member.

clear hardware stack-unit (for high-Gigabit Ethernet ports)

Clear statistics from selected hardware components.

NOTE: Only the parameters that are newly introduced with this command in Release 9.3(0.0) are explained here. For a complete description of all of the options that are available with this command, refer the relevant *Command Reference Guide* of the applicable platform of Release 9.2(0.0).

Z9000

Syntax

```
clear hardware stack-unit number {counters | unit number
counters | hg-stats [unit number [port port-number]] | cpu
data-plane statistics | cpu i2c statistics | stack-port number}
```

Parameters

hg-stats [unit *unit-number* [port *port-number* | no]]	Enter the keyword `hg-stats` to display high–Gigabit Ethernet or backplane port buffer and queue statistics on the selected stack member. Optionally, use the keyword `unit` with a number to select port-pipe 0 to 5, and then use `port` *`port-number`* to select a port on that port-pipe. For Z9000, valid backplane ports for leaf NPU units (units 0–3) range from 34-41 and for spine NPU units (units 4–5) range from 1 -16.

Defaults none

Command Modes EXEC Privilege

Command History

Version 9.3.0.0	Added support for the `hg-stats` option on the Z9000 platform.

Related Commands show hardware stack-unit — displays the data plane or management plane input and output statistics of the designated component of the designated stack member.

Viewing Buffer Utilization and Queue Statistics on High-Gigabit Ethernet Backplane Ports

You can view the buffer utilization and queue statistical counters for high-Gigabit Ethernet ports or trunk groups that operate as backplane ports. This functionality is supported on the Z9000 platform.

You can now view the queue statistics and buffer utilization counters for the internal leaf port and spine port queues on a Z9000 platform using the appropriate show commands. Transmit and receive counters, and drop counters per queue are computed and displayed for internal queues on high-Gigabit Ethernet ports in the leaf and spine nodes on the Z9000 platform. Buffer utilization counters supported for front-end ports are extended to high-Gigabit backplane ports.

Buffer counters include a new metric or parameter called Total Count cells. This field contains the total number of cells currently being used by all queues on all ports in a portpipe. The f10-bp-stats.mib is for statistics collection of backplane ports. For Z9000, valid backplane ports for leaf NPU units range from 34-41 and for spine NPU units range from 1-16. In a Card Type (slot), NPUT units are always indexed starting with the leaf NPU units and then proceeding to the spine NPU units.

In an NPU unit, the port numbering of backplane local ports starts from the end of the last front-end local port ID used.

Until Dell Networking OS Release 9.2.0.0, the show commands display the statistics other than the details computed by the buffer statistics tracking counters for the egress queues. You can now use the relevant show commands to display the ingress counters that are not part of the counters that are calculated by the buffer statistics tracking method for each port per priority group. You can use the `show hardware stack-unit <unit-num> buffer unit <unit-num>` command to display the buffer statistics and queue information. You can use the `clear hardware stack-unit <unit-num>` command to reset the statistical details associated with high-Gigabit Ethernet ports. This functionality is supported on the Z9000 platform.

The following commands are enhanced to display the buffer statistics tracking counters for high-Gigabit backplane ports on the Z9000 platform:

- `show hardware stack-unit 0 buffer unit 0 total-buffer`

  ```
  -----  Buffer Details for Unit 0 -----
   Used Packet Buffer for the Unit: 0
   Current Available Packet Buffer for the Unit: 46080
   Is Dynamic Packet Buffer allocate for the unit: TRUE
  FTOS#
  ```

  In the preceding sample output that is displayed, which is a portion of the complete output that is shown when you run this command, the shared buffer space that is available to be allotted to the specific port for the corresponding stack unit, the shared buffer space that is in use by the packets, and whether dynamic packet buffering allocation is activated are displayed.

- `show hardware stack-unit 0 buffer unit {0-5} port all buffer-info`—Supports backplane or high-Gigabit ports of all queues for all ports in a specific unit.
- `show hardware stack-unit 0 buffer unit {0-5} port {1-41} queue {1-14} buffer-info`—Supports backplane HG ports of all units for a specific port and queue in a unit.
- `show hardware stack-unit 0 buffer unit {0-5} port {1-41} buffer-info`—Supports backplane high-Gigabit port for switch fabric or spine units for a specific port.
- `show hardware stack-unit 0 buffer unit {0-5} port all queue all buffer-info`—Supports backplane high-Gigabit ports of all units for all queues of all ports in a unit.

- `show hardware stack-unit 0 buffer unit {0-5} port {1-41} queue all buffer-info`—Supports backplane high-Gigabit port for switch fabric or spine units for a specific port and all queues.
- `show hardware stack-unit 0 drops unit {0-5} port {1-41}`—Supports drop counters for non-fanout high-Gigabit ports (backplane ports).

12

Miscellaneous Settings

This chapter contains several, diversified behavioral-changes and enhancements that apply to this release.

- Default Host Name Change
- hostname (for Changes to Default)

Setting a Threshold for Switching to the SPT

The functionality to specify a threshold for switchover to the shortest path trees (SPTs) is available on S-Series platforms. After a receiver receives traffic from the RP, PM-SM switches to SPT to forward multicast traffic. Every multicast group has an RP and a unidirectional shared tree (group-specific shared tree).

The SPT-Threshold is zero, which means that the last-hop designated router (DR) joins the shortest path tree (SPT) to the source upon receiving the first multicast packet.

Initially, a single PIM-SM tree called a shared tree to distribute traffic. It is called shared because all traffic for the group, regardless of the source, or the location of the source, must pass through the RP. The shared tree is unidirectional; that is, all multicast traffic flows only from the RP to the receivers. Once a receiver receives traffic from the RP, PM-SM switches to *shortest path trees* (SPT) to forward multicast traffic, which connects the receiver directly to the source.

You can configure PIM to switch over to the SPT when the router receives multicast packets at or beyond a specified rate.

Table 4. Configuring PIM to Switch Over to the SPT

IPv4	Configure PIM to switch over to the SPT when the multicast packet rate is at or beyond a specified rate. The keyword infinity directs PIM to never switch to the SPT.	ip pim spt-threshold {value | infinity} Default: 10 kbps	CONFIGURATION
IPv6	Configure PIM to switch over to the SPT when the multicast packet rate is at or beyond a specified rate. The keyword infinity directs PIM to never switch to the SPT.	ip pim spt-threshold {value | infinity} Default: 10 kbps	CONFIGURATION

ip pim spt-threshold

To switch to the shortest path tree when the traffic reaches the specified threshold value, configure the PIM router.

S6000

Syntax `ip pim spt-threshold value | infinity`

To return to the default value, use the `no ip pim spt-threshold` command.

Parameters	*value*	(OPTIONAL) Enter the traffic value in kilobits per second. The default is **10 packets per second**. A value of zero (0) causes a switchover on the first packet.
	infinity	(OPTIONAL) Enter the keyword `infinity` to never switch to the source-tree.
Defaults	Not configured.	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the S6000.
Usage Information	This command is applicable to last hop routers on the shared tree towards the rendezvous point (RP).	

ip route bfd (for S6000)

Enable BFD for all neighbors configured through static routes.

S6000

Syntax

`ip route bfd [interval` *`interval`* `min_rx` *`min_rx`* `multiplier` *`value`* `role {active | passive}]`

To disable BFD for all neighbors configured through static routes, use the `no ip route bfd [interval` *`interval`* `min_rx` *`min_rx`* `multiplier` *`value`* `role {active | passive}]` command.

Parameters

interval *milliseconds*	(OPTIONAL) Enter the keywords `interval` to specify non-default BFD session parameters beginning with the transmission interval. The range is from 50 to 1000. The default is **100**.
min_rx *milliseconds*	Enter the keywords `min_rx` to specify the minimum rate at which the local system receives control packets from the remote system. The range is from 50 to 1000. The default is **100**.
multiplier *value*	Enter the keywords `multiplier` to specify the number of packets that must be missed in order to declare a session down. The range is from 3 to 50. The default is **3**.
role [active | passive]	Enter the role that the local system assumes: • `Active` — The active system initiates the BFD session. Both systems can be active for the same session. • `Passive` — The passive system does not initiate a session. It only responds to a request for session initialization from the active system. The default is **Active**.

Defaults	See Parameters	
Command Modes	CONFIGURATION	
Command History	**Version 9.3. (0.0)**	Introduced on S6000.
Related Commands	show bfd neighbors – displays the BFD neighbor information on all interfaces or a specified interface.	

Configure BFD for Static Routes

Configuring BFD for static routes is supported on Z9000 S4810 S4820T S6000.

BFD offers systems a link state detection mechanism for static routes. With BFD, systems are notified to remove static routes from the routing table as soon as the link state change occurs, rather than waiting until packets fail to reach their next hop.

Configuring BFD for static routes is a three-step process:

1. Enable BFD globally.
2. Configure static routes on both routers on the system (either local or remote).
3. Configure an IP route to connect BFD on the static routes using the `ip route bfd` command.

Related Configuration Tasks

- Changing Static Route Session Parameters
- Disabling BFD for Static Routes

Changing Static Route Session Parameters

BFD sessions are configured with default intervals and a default role.

The parameters you can configure are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured for all static routes. If you change a parameter, the change affects all sessions for static routes.

To change parameters for static route sessions, use the following command .

- Change parameters for all static route sessions.

  CONFIGURATION mode

  ```
  ip route bfd interval milliseconds min_rx milliseconds multiplier value role
  [active | passive]
  ```

To view session parameters, use the `show bfd neighbors detail` command, as shown in the examples in Displaying BFD for BGP Information.

Establishing Sessions for Static Routes

Sessions are established for all neighbors that are the next hop of a static route.

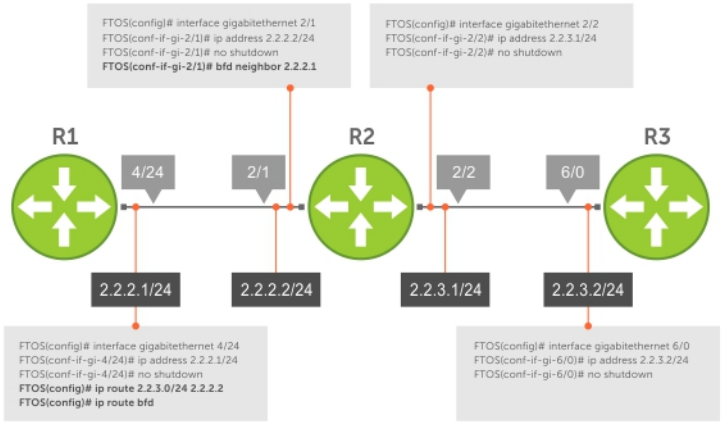

Figure 1. Establishing Sessions for Static Routes

To establish a BFD session, use the following command.

- Establish BFD sessions for all neighbors that are the next hop of a static route.
  CONFIGURATION mode

  `ip route bfd`

To verify that sessions have been created for static routes, use the `show bfd neighbors` command.

Example of the `show bfd neighbors` Command to Verify Static Routes

```
R1(conf)#ip route 2.2.3.0/24 2.2.2.2
R1(conf)#ip route bfd
R1(conf)#do show bfd neighbors

* - Active session role
Ad Dn - Admin Down
C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)
LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
2.2.2.1   2.2.2.2    Gi 4/24   Up    100    100    4    R
```

To view detailed session information, use the `show bfd neighbors detail` command, as shown in the examples in Displaying BFD for BGP Information.

Disabling BFD for Static Routes

If you disable BFD, all static route BFD sessions are torn down.

A final Admin Down packet is sent to all neighbors on the remote systems, and those neighbors change to the Down state.

To disable BFD for static routes, use the following command.

- Disable BFD for static routes.
  CONFIGURATION mode
  `no ip route bfd`

source (port monitoring for 40-Gigabit Ethernet)

Configure a port monitor source and destination. Starting with Dell Networking OS Release 9.3(0.0), you can also configure a 40-Gigabit Ethernet interface as the destination interface or port to which the monitored traffic is sent .

Syntax

`source` *`interface`* `destination` *`interface`* `direction {rx | tx | both}`

To disable a monitor source, use the `no source` *`interface`* `destination` *`interface`* `direction {rx | tx | both}` command.

Parameters

interface	Enter the one of the following keywords and slot/port information: • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
destination	Enter the keyword `destination` to indicate the interface destination.
direction {rx | tx | both}	Enter the keyword `direction` then one of the packet directional indicators. • `rx`: to monitor receiving packets only. • `tx`: to monitor transmitting packets only. • `both`: to monitor both transmitting and receiving packets.

Defaults none

Command Modes MONITOR SESSION (conf-mon-*sess-session-ID*)

Command History

Version 9.3(0.0)	Added support for the `fortyGigE` keyword on M I/O Aggregator
Version 8.3.17.0	Supported on M I/O Aggregator

Example

```
Dell(conf-mon-sess-11)#source fortygi 10/0 destination gi
10/47 direction rx
Dell(conf-mon-sess-11)#
```

13

Microsoft Network Load Balancing

This functionality is supported on the S4810, S4820T, S6000, Z9000, and MXL platforms.

Network Load Balancing (NLB) is a clustering mechanism that is implemented by Microsoft on Windows 2000 Server and Windows Server 2003 operating systems. NLB uses a distributed methodology or pattern to equally split and balance the network traffic load across a set of servers that are part of the cluster or group. NLB combines the servers into a single multicast group and attempts to use the standard multicast IP or unicast IP addresses, and MAC addresses for transmission of network traffic. At the same time, it also uses a single virtual IP address for all clients as the destination IP address, which enables servers to join the same multicast group in a way that is transparent to the clients (the clients do not notice the addition of new servers to the group). The clients use a cluster IP address to connect to the server. The NLB functionality enables flooding of traffic over the VLAN ports (for unicast mode) or a subset of ports in a VLAN (for multicast mode) to avoid overloading and effective performance of the servers for optimal processing of data packets.

NLB functions in two modes, namely, unicast mode and multicast mode. The cluster IP address and the associated cluster MAC address are configured in the NLB application running on the Windows Server. In unicast mode, when the server IP address is attempted to be resolved to the MAC address using the ARP application, the switch determines the ARP reply to be an NLB type of ARP reply obtained from the server. The switch then maps the IP address (cluster IP) with the MAC address (cluster MAC address). In multicast mode, the cluster IP address is mapped to a cluster multicast MAC address that is configured using a static ARP CLI configuration command. After the NLB entry is learned, the traffic is forwarded to all the servers in the VLAN corresponding to the cluster virtual IP address.

NLB Unicast Mode Scenario

Consider a sample topology in which four servers, namely S1 through S4, are configured as a cluster or a farm. This set of servers is connected to a Layer 3 switch, which in turn is connected to the end-clients. The servers contain a single IP address (IP-cluster address of 172.16.2.20) and a single unicast MAC address (MAC-Cluster address of 00-bf-ac-10-00-01) for load-balancing. Because multiple ports of a switch cannot learn a single MAC address, the servers are assigned with MAC addresseses of MAC-s1 to MAC-s4) respectively on S1 through S4 in addition to the MAC cluster address. All the servers of the cluster belong to the VLAN named VLAN1.

In unicast NLB mode, the following sequence of events occurs:

- The switch sends an ARP request to resolve the IP address to the cluster MAC address.
- The ARP servers send an ARP response with the MAC cluster address in the ARP header and a MAC address of MAC-s1/s2/s3/s4 (for servers S1 through S4) in the Ethernet header.
- The switch associates the IP address with the MAC cluster address with the last ARP response it obtains. Assume that in this case, the last ARP reply is obtained from MAC-s4.(assuming that the ARP response with MAC-s4 is received as the last one). The interface associated with server, S4, is added to the ARP table.
- With NLB feature enabled, after learning the NLB ARP entry, all the subsequent traffic is flooded on all ports in VLAN1.

With NLB, the data frame is forwarded to all the servers for them to perform load-balancing.

NLB Multicast Mode Scenario

Consider a sample topology in which four servers, namely S1 through S4, are configured as a cluster or a farm. This set of servers is connected to a Layer 3 switch, which in turn is connected to the end-clients. They contain a single multicast MAC address (MAC-Cluster: 03-00-5E-11-11-11).

In multicast NLB mode, a static ARP configuration command is configured to associate the cluster IP address with a multicast cluster MAC address.

With multicast NLB mode, the data is forwarded to all the servers based on the port specified the Layer 2 multicast command, which is the `mac-address-table static <multicast_mac> multicast vlan <vlan_id> output-range <port1>, <port2>` command in CONFIGURATION mode.

Limitations With Enabling NLB on Switches

The following limitations apply to switches on which you configure NLB:

- The NLB unicast mode uses switch flooding to transmit all packets to all the servers that are part of the VLAN. When a large volume of traffic is processed, the clustering performance might be impacted in a small way. This limitation is applicable to switches that perform unicast flooding in the software.
- The `ip vlan-flooding` command applies globally across the system and for all VLANs. In cases where the ARP replies contain a discrepancy with the Ethernet SHA and ARP header SHA frames and NLB is applicable, flooding of packets over the relevant VLAN occurs.
- The maximum number of concurrent clusters that is supported is eight.

Benefits and Working of Microsoft Clustering

Microsoft Clustering allows multiple servers using Microsoft Windows to be represented by one MAC address and IP address in order to provide transparent failover or balancing. FTOS does not recognize server clusters by default; it must be configured to do so. When an ARP request is sent to a server cluster, either the active server or all of the servers send a reply, depending on the cluster configuration. If the active server sends a reply, the Dell Force10 switch learns the active server's MAC address. If all servers reply, the switch registers only the last received ARP reply, and the switch learns one server's actual MAC address; the virtual MAC address is never learned. Because the virtual MAC address is never learned, traffic is forwarded to only one server rather than the entire cluster, and failover and balancing are not preserved.

To preserve failover and balancing, the switch forwards the traffic destined for the server cluster out all member ports in the VLAN connected to the cluster. To ensure that this happens, you must configure the command `ip vlan-flooding` on the Dell Force10 switch at the time that the Microsoft cluster is configured. The server MAC address is given in the Ethernet frame header of the ARP reply, while the virtual MAC address representing the cluster is given in the payload. Then, all traffic destined for the cluster is flooded out of all member ports. Since all of the servers in the cluster receive traffic, failover and balancing are preserved.

Enable and Disable VLAN Flooding

- The older ARP entries are overwritten whenever newer NLB entries are learned.
- All ARP entries learned after the feature is enabled are deleted when the feature is disabled, and RP2 triggers ARP resolution. The feature is disabled with the command `no ip vlan-flooding`.

- When a port is added to the VLAN, the port automatically receives traffic if the feature is enabled. Old ARP entries are not deleted or updated.
- When a member port is deleted, its ARP entries are also deleted from the CAM.
- Port channels in the VLAN also receive traffic.
- There is no impact on the configuration from saving the configuration.
- The feature if enabled is displayed in the `show running-config` command output that displays `ip vlan-flooding` CLI configuration. Apart from it, there is no indication of the enabling of this capability.

Configuring a Switch for NLB

This functionality is supported on the S4810, S4820T, S6000, Z9000, and MXL platforms.

To enable a switch for unicast NLB mode of functioning, perform the following:

Specify that all the Layer 3 unicast routed data traffic going through a VLAN member port need to be flooded across all the member ports of that VLAN by entering the `ip vlan-flooding` command. There might be some ARP table entries which are resolved through ARP packets which had Ethernet MAC SA different from MAC information inside the ARP packet. This unicast data traffic flooding occurs only for those packets which use these ARP entries.

CONFIGURATION mode

```
ip vlan-flooding
```

To enable a switch for multicast NLB mode of functioning, perform the following:

1. For multicast mode of NLB, to associate an IP address with a multicast MAC address in the switch, add a static ARP entry by entering the `arp ip-address multicast-mac-address` command in Global configuration mode. This setting causes the multicast MAC address to be mapped to the cluster IP address for NLB mode of operation of the switch.

   INTERFACE mode

   ```
   arp ip-address multicast-mac-address interface
   ```
2. Associate specific MAC or hardware addresses to VLANs.

   CONFIGURATION mode

   ```
   mac-address-table static multicast-mac-address vlan vlan-id output-range interface
   ```

arp (for Multicast MAC Address)

To associate an IP address with a multicast MAC address in the switch when you configure multicast mode of network load balancing (NLB), use address resolution protocol (ARP).

Syntax

```
arp ip-address multicast-mac-address interface
```

To remove an ARP address, use the `no arp ip-address` command.

Parameters

ip-address	Enter an IP address in dotted decimal format.

multicast-mac-address	Enter a 48-bit hexadecimal address in nn:nn:nn:nn:nn:nn format for the static MAC address to be used to switch multicast traffic.
interface	Enter any of the following keywords and slot/port or number information: • For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information. • The interface specified here must be one of the interfaces configured using the `{output-range | output} interface` option with the `mac-address-table static` command.

Defaults	Not configured.	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Added support for association of an IP address with multicast MAC address on the S4810, S4820T, S6000, Z9000, and MXL platforms.
Usage Information	For multicast mode of NLB, to associate an IP address with a multicast MAC address in the switch, use address resolution protocol (ARP) by entering the `arp ip-address multicast-mac-address` command in Global configuration mode. This setting causes the multicast MAC address to be mapped to the cluster IP address for NLB mode of operation of the switch.	
Related Commands	clear arp-cache — clears dynamic ARP entries from the ARP table. show arp — displays the ARP table.	

mac-address-table static (for Multicast MAC Address)

For multicast mode of network load balancing (NLB), configure a static multicast MAC address, associate the multicast MAC address with the VLAN used to switch Layer 2 multicast traffic, and add output ports that will receive multicast streams on the VLAN. To delete a configured static multicast MAC address from the MAC address table on the router, enter the `no mac-address-table static multicast-mac-address` command.

Syntax

```
mac-address-table static multicast-mac-address multicast vlan
vlan-id range-output {single-interface | interface-list |
interface-range}
```

To remove a MAC address, use the `no mac-address-table static multicast-mac-address output interface vlan vlan-id` command.

Parameters

multicast-mac-address	Enter the 48-bit hexadecimal address in nn:nn:nn:nn:nn:nn format.
multicast	Enter a vlan port to where L2 multicast MAC traffic is forwarded. **NOTE:** Use this option if you want multicast functionality in an L2 VLAN without IGMP protocols.
output *interface*	For a multicast MAC address, enter the keyword `output` then one of the following interfaces for which traffic is forwarded: • For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
output-range *interface*	For a multicast MAC address, enter the keyword `output-range` then one of the following interfaces to indicate a range of ports for which traffic is forwarded: • For a 1-Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information. • For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
vlan *vlan-id*	Enter the keyword `vlan` then a VLAN ID number from 1 to 4094.

Defaults

Not configured.

Command Modes

CONFIGURATION

Command History

The following is a list of the FTOS version history for this command.

Version 9.3.0.0	Added support for multicast MAC address on the S4810, S4820T, S6000, Z9000, and MXL platforms.

Usage Information

When a multicast source and multicast receivers are in the same VLAN, you can configure a router so that multicast traffic is switched only to the ports assigned to a VLAN that is associated with a static multicast MAC address. However, before you can configure a static MAC address and associate it with a VLAN used to switch Layer 2 multicast traffic, you must first enable the router for Layer 2 multicast switching with the `ip multicast-mode l2` command.

Example (Multicast)	`mac-address-table static 01:00:5E:01:00:01 {multicast vlan 2 output—range Te 0/2,Te 0/3}`

ip vlan-flooding

Enable unicast data traffic flooding on VLAN member ports.

Syntax	`ip vlan-flooding` To disable, use the `no ip vlan-flooding` command.
Command Modes	CONFIGURATION
Command History	**Version 9.3.0.0** Introduced on the S4810, S4820T, S6000, Z9000, and MXL platforms
Default	Disabled
Usage Information	By default this command is disabled. There might be some ARP table entries which are resolved through ARP packets which had Ethernet MAC SA different from MAC information inside the ARP packet. This unicast data traffic flooding occurs only for those packets which use these ARP entries.

14

Quality of Service (QoS)

This chapter describes the QoS enhancements and contains the following sections:

- Classifying Packets Based on a Combination of DSCP Code Points and VLAN IDs
- Specifying Policy-Based Rate Shaping in Packets Per Second
- Managing Hardware Buffer Statistics
- Classifying Layer 2 Traffic on Layer 3 Interfaces
- RRoCE Overview

Specifying Policy-Based Rate Shaping in Packets Per Second

The capability to configure rate shaping for QoS output policies in packets per second (pps) is supported on the S6000 platform.

You can configure rate shaping that is applied to a QoS output policy in packets per second (pps), apart from specifying the rate shaping value in bytes. You can also configure the peak rate, which is the maximum permissible rate for the packets, and the committed rate, which is the minimum confirmed rate that is maintained for the packets, in kilobits per second (Kbps) or pps.

Committed rate refers to the guaranteed bandwidth for traffic entering or leaving the interface under normal network conditions. When traffic propagates at an average rate that is less than or equal to the committed rate, it is considered to be green-colored or coded. When the transmitted traffic falls below the committed rate, the bandwidth that is not used by any traffic that is traversing the network is aggregated to form the committed burst size. Traffic is considered to be green-colored up to a point at which the unutilized bandwidth does not exceeded the configured committed burst size.

Peak rate refers to the maximum rate for traffic arriving or exiting an interface under normal traffic conditions. Peak burst size indicates the maximum size of unused peak bandwidth that is aggregated. This aggregated bandwidth enables brief durations of burst traffic that exceeds the peak rate and committed burst.

In releases of Dell Networking OS earlier than Release 9.3.0.0, you can configure only the maximum shaping attributes, such as the peak rate and peak burst settings. You can now specify the committed or minimum burst and committed rate attributes. The committed burst and committed rate values can be defined either in bytes or packets per second.

You can use the `rate-shape pps peak-rate` *`burst-packets`* command in the QoS Policy Out Configuration mode to configure the peak rate and burst size as a measure of pps. Alternatively, you can use the `rate shape kbps peak-rate` *`burst-KB`* command to configure the peak rate and peak burst size as a measure of bytes.

Similarly, you can use the `rate-shape pps peak-rate` *`burst-packets`* `committed pps committed-rate` *`burst-packets`* command in the QoS Policy Out Configuration mode to configure the committed rate and committed burst size as a measure of pps. Alternatively, you can use the `rate shape kbps peak-rate` *`burst-KB`* `committed kbps committed-rate` *`burst-KB`* command to configure the committed rate and committed burst size as a measure of bytes. If you configure the peak

rate in pps, the peak burst size must also be configured as a measure of number of packets. Similarly, if you configure the peak rate in Kbps, the peak burst size must also be configured as a measure of bytes.

Configuring Policy-Based Rate Shaping

The capability to configure rate shaping for QoS output policies in packets per second (pps) is supported on the S6000 platform.

You can explicitly specify the rate shaping functionality for QoS output policies as peak rate and committed rate attributes. You can also configure the peak burst and committed burst sizes. All of these settings can be configured in Kbps, Mbps, or pps.

To configure the peak and committed rates and burst sizes, perform the following:

1. Configure the peak rate and peak burst size in pps in QoS Policy Out Configuration mode.
   QOS-POLICY-OUT mode
   ```
   Dell(config-qos-policy-out)# rate shape pps peak-rate burst-packets
   ```
2. Alternatively, configure the peak rate and peak burst in bytes.
   QOS-POLICY-OUT mode
   ```
   Dell(config-qos-policy-out)# rate shape Kbps peak-rate burst-KB
   ```
3. Configure the committed rate and committed burst size in pps.
   QOS-POLICY-OUT mode
   ```
   Dell(config-qos-policy-out)# rate shape pps peak-rate burst-packets
   committed pps committed-rate burst-packets
   ```
4. Alternatively, configure the committed rate and committed burst in bytes.
   QOS-POLICY-OUT mode
   ```
   Dell(config-qos-policy-out)# rate shape Kbps peak-rate burst-KB committed
   Kbps committed-rate burst-KB
   ```

Configuring Weights and ECN for WRED

The mechanism to configure a weight factor for WRED and ECN functionality for backplane ports is supported on the S6000 platform. However, this mechanism to configure a weight for WRD and ECN functionality for front-end ports is supported on the S6000 and Z9000 platforms.

Weighted random early detection (WRED) congestion avoidance mechanism drops packets to prevent buffering resources from being consumed. Traffic is a mixture of various kinds of packets. The rate at which some types of packets arrive might be greater than others. In this case, the space on the buffer and traffic manager (BTM) (ingress or egress) can be consumed by only one 656 or a few types of traffic, leaving no space for other types. You can apply a WRED profile to a policy-map so that specified traffic can be prevented from consuming too much of the BTM resources.

WRED drops packets when the average queue length exceeds the configured threshold value to signify congestion. Explicit Congestion Notification (ECN) is a capability that enhances WRED by marking the packets instead of causing WRED to drop them when the threshold value is exceeded. If you configure ECN for WRED, devices employ this functionality of ECN to mark the packets and reduce the rate of sending packets in a congested, heavily-loaded network.

In a best-effort network topology, data packets are transmitted in a manner in which latency or throughput are not maintained to be at the effective, most-optimal level. Packets are dropped when the network experiences a large traffic load. This best-effort network deployment is not suitable for applications that are time-sensitive such as video on demand (VoD) or voice over IP (VoIP) applications. In

such cases, you can use ECN in conjunction with WRED to resolve this problem of dropping of packets under congested conditions.

Using ECN, instead of entirely dropping the packets when the network experiences excessive traffic, the packets are marked for transmission at a later time after the network recovers from the heavy traffic state to optimal load. In this manner, enhanced performance and throughput is achieved. Also, the devices can respond to congestion before a queue overflows and packets are dropped, enabling improved queue management.

When a packet reaches the device with ECN enabled for WRED, the average queue size is computed. To measure the average queue size, a weight factor is used. This weight factor is user-configurable. You can use the wred weight number command to configure the weight for the WRED average queue size. If the average queue size is less than the minimum threshold of WRED, the received packet is queued. If the average queue size is more than the maximum threshold of WRED, the packet is dropped. If the average queue size is between the minimum and maximum threshold values, the decision to drop or queue the packet is taken based on the packet drop probability. The probability that a packet is dropped depends on the minimum threshold, maximum threshold, and mark probability denominator. The rate of packet drop increases in a proportional way as the average queue size increases, until the average queue size reaches the maximum threshold value. The mark probability value is the number of packets dropped when the average queue size reaches the maximum threshold value.

The average queue size is computed using the preceding average size and the current queue size. The following is the formula to calculate the average queue size: average-queue-size (t+1) = average-queue-size (t) + (current-queue-length - average-queue-size (t))/2^N

where t is the time or the current instant at which average queue size is measured, t+1 is the next time iteration at which average queue size is calculated, and N is the weight factor.

The weight factor is set to zero by default, which causes the same behavior as dropping of packets by WRED during network loads or also called instantaneous ECN marking. In a topology in which congestion of the network varies over time, you can specify a weight to enable a smooth, seamless averaging of packets to handle bursty nature of the packets based on the previous time sampling performed. You can specify the weight parameter for front-end and backplane ports separately in the range of 0 through 15.

You can enable WRED and ECN capabilities per queue for granularity. You can disable these functionalities per queue, and you can also specify the minimum and maximum buffer thresholds for each color-coding of the packets. You can configure maximum drop rate percentage or yellow and green profiles. You can set up these parameters for both front-end and backplane ports.

Global Service Pools With WRED and ECN Settings

A global buffer pool that is a shared buffer pool accessed by multiple queues when the minimum guaranteed buffers for the queue are consumed can be configured on the S6000 and Z9000 platforms.

Support for global service pools is now available. You can configure global service pools that are shared buffer pools accessed by multiple queues when the minimum guaranteed buffers for the queue are consumed. S4810, S4820T, S6000, and Z9000 platforms support four global service-pools in the egress direction. Two service pools are used– one for lossy queues and the other for lossless (priority-based flow control (PFC)) queues. You can enable WRED and ECN configuration on the global service-pools.

You can define WRED profiles and weight on each of the global service-pools for both lossy and lossless (PFC) service- pools. The following events occur when you configure WRED and ECN on global service pools:

- If WRED/ECN is enabled on the global service-pool with threshold values and if it is not enabled on the queues, WRED/ECN are not effective based on global service-pool WRED thresholds. The queue on which traffic is scheduled must contain WRED/ECN settings enabled for WRED to be valid for that traffic.

- When WRED is configured on the global service-pool (regardless of whether ECN on global service-pool is configured), and one or more queues have WRED enabled and ECN disabled, WRED is effective for the minimum of the thresholds between the queue thresholed and the service pool threshold.
- When WRED is configured on the global service-pool (regardless of whether ECN on global service-pool is configured), and one or more queues are enabled with both WRED and ECN, ECN marking takes effect. The packets are ECN marked up to shared- buffer limits as determined by the shared-ratio for that global service-pool.

WRED/ECN configurations for the queues that belong to backplane ports are common to all the backplane ports and cannot be specified separately for each backplane port granularity. This behavior occurs to prevent system-level complexities in enabling this support for backplane ports. Also, WRED/ECN is not supported for multicast packets.

The following table describes the WRED and ECN operations that occur for various scenarios of WRED and ECN configuration on the queue and service pool. (X denotes not-applicable in the table, 1 indicates that the setting is enabled, 0 represents a disabled setting.)

Table 5. Scenarios of WRED and ECN Configuration

Queue Configuration		Service-Pool Configuration		WRED Threshold Relationship Q threshold = Q-T, Service pool threshold = SP-T	Expected Functionality
WRED	ECN	WRED	ECN		
0	0	X	X	X	WRED/ECN not applicable
1	0	0	X	X	Queue based WRED, No ECN marking
		1	X	Q-T < SP-T	
				SP-T < Q-T	SP based WRED, No ECN marking
1	1	0	X	X	Queue based ECN marking above queue threshold. ECN marking up to shared buffer limits of the service-pool and then packets are tail dropped.
		1	X	Q-T < SP-T	
				SP-T < Q-T	Same as above but ECN marking starts above SP-T.

Configuring WRED and ECN Attributes

The mechanism to configure a weight factor for WRED and ECN functionality for backplane ports is supported on the S6000 platform. However, this mechanism to configure a weight for WRED and ECN functionality for front-end ports is supported on the S6000 and Z9000 platforms. A global buffer pool that is a shared buffer pool accessed by multiple queues when the minimum guaranteed buffers for the queue are consumed can be configured on the S6000 and Z9000 platforms.

WRED drops packets when the average queue length exceeds the configured threshold value to signify congestion. Explicit Congestion Notification (ECN) is a capability that enhances WRED by marking the packets instead of causing WRED to drop them when the threshold value is exceeded. If you configure ECN for WRED, devices employ this functionality of ECN to mark the packets and reduce the rate of sending packets in a congested, heavily-loaded network.

To configure the weight factor for WRED and ECN capabilities, global buffer pools for multiple queues, and associating a service class with ECN marking, perform the following:

1. Configure the weight factor for computation of average-queue size. This weight value applies to front-end ports.

   QOS-POLICY-OUT mode

   ```
   Dell(conf-qos-policy-out)#wred weight number
   ```

2. Configure a WRED profile, and specify the threshold and maximum drop rate

   WRED mode

   ```
   Dell(conf-wred) #wred thresh-1
   Dell(conf-wred) #threshold min 100 max 200 max-drop-rate 40
   ```

3. Configure another WRED profile, and specify the threshold and maximum drop rate

   WRED mode

   ```
   Dell(conf-wred) #wred thresh-2
   Dell(conf-wred) #threshold min 300 max 400 max-drop-rate 80
   ```

4. Associate the service class with the WRED profile, and assign the WRED profile to specific queues on backplane ports

   CONFIGURATION mode

   ```
   Dell(conf) #service-class wred green queue5 thresh-1 queue7 thresh-2
   backplane
   Dell(conf) #service-class wred yellow queue1 thresh-2 queue3 thresh-1
   backplane
   Dell(conf) #service-class wred weight queue0 11 queue6 4 queue7 9 backplane
   ```

5. Create a global buffer pool that is a shared buffer pool accessed by multiple queues when the minimum guaranteed buffers for the queue are consumed. S4810, S4820T, and S6000 platforms support four global service-pools in the egress direction. The Z9000 platform supports only pool 0.

   mode

   ```
   Dell(conf) #service-pool wred green pool0 thresh-1 pool1 thresh-2
   Dell(conf) #service-pool wred yellow pool0 thresh-3 pool1 thresh-4
   Dell(conf) #service-pool wred weight pool0 11 pool1 4
   ```

6. Attach the ECN marking to specific queues on backplane ports with a service class

   CONFIGURATION mode

   ```
   Dell(conf) #service-class wred ecn 0, 3-5, 7 backplane
   ```

7. Create a service class and associate the threshold weight of the shared buffer with each of the queues per port in the egress direction.

   INTERFACE mode

   ```
   Dell(conf-if-te-0/8)#Service-class buffer shared-threshold-weight queue5 4
   queue7 6
   ```

Classifying Layer 2 Traffic on Layer 3 Interfaces

You can configure VLAN tags on a physical Layer 3 interface (that is configured with an IP address and is not associated with any VLAN) to enable Layer 3 packets that contain Dot1p—(IEEE 802.1p) Packet classification (Layer 2 headers) to be processed properly. You can thereby enable classification of Layer 2 packets on L3 interfaces (ports that are not configured as switch ports). You can configure a VLAN subinterface over a physical underlying interface and classify packets using the dot1p value.

You can use the service-policy input policy-name layer 2 command in Interface Configuration mode to apply an input policy map to Layer 3 physical interfaces.

To apply a Layer 2 policy on Layer 3 interfaces, perform the following:

1. Configure an interface with an IP address or a VLAN subinterface
   CONFIGURATION mode
   ```
   Dell(conf)# int fo 0/0
   ```
   INTERFACE mode
   ```
   Dell(conf-if-fo-0/0)# ip address 90.1.1.1/16
   ```
2. Configure Layer2 policy with Layer 2 (Dot1p or source MAC-based) classification rules.
   CONFIGURATION mode
   ```
   Dell(conf)# policy-map-input l2p layer2
   ```
3. Apply the L2 policy on the Layer 3 interface.
   INTERFACE mode
   ```
   Dell(conf-if-fo-0/0)# service-policy input l2p layer2
   ```

Managing Hardware Buffer Statistics

Bufffer statistics tracking utility is supported on the S6000 platform.

The memory management unit (MMU) on S6000 and Z9000 platforms is 12.2 MB in size. It contains approximately 60,000 cells, each of which is 208 bytes in size. MMU also has another portion of 3 MB allocated to it. The entire MMU space is shared across a maximum of 104 logical ports to support egress admission-control mechanisms to implement scheduling and shaping on per-port and per-queue levels. Also, the MMU buffer cells can be used by each port or queue either as a static partition or as a dynamic partition. With dynamic mode, you can specify the percentage of available buffer that is utilized by a queue. This dynamic partition or block is set to be two-thirds of the available buffers for all unicast queues and one-fifth of the available buffers for all multicast queues on these platforms.

The maximum number of ports, including fan-out, supported is 64 and the maximum number of queues supported is 15. Analyzing and evaluating buffer statistics is a mechanism that enables monitoring of resources and tuning of allocation of buffers. This mechanism operates in two modes, namely, Max Use count mode and Current Use count mode. Max Use count mode provides the maximum values of counters accumulated over a period of time. Current Use count mode enables you to obtain a snapshot of the counters at a particular point in me using a triggering utility. The trigger can either be software-based or based on a predetermined threshold event. Software-based triggers are supported, which are the values derived from the show command output in the Max use count mode. In Dell Networking OS Release 9.3.0.0, only the max use count mode of operation is supported for computation of maximum counter values.

Depending on the buffer space statistical values that you can obtain, you can modify the settings for buffer area to achieve enhanced reliability and efficiency in the handling of packets. This evaluation and administration of buffer statistics is useful and important in deployments that experience congestion frequently. The receive buffer must be large enough to save all data that is received when the system processes a PFC PAUSE frame.

You can use the `service-class buffer shared-threshold-weight queue0 ... queue7` *`number`* command in Interface Configuration mode to specify the threshold weight for the shared buffer for each of the queues per port.

1. Create a 10-Gigabit Ethernet interface.
   ```
   Dell(conf)#interface TenGigabitEthernet 0/8
   ```

2. Configure the threshold weight of the shared buffer for the queues you want. In this example, this setting is configured for queues 5 and 7.

   ```
   Dell(conf-if-te-0/8)#Service-class buffer shared-threshold-weight queue5 4
   queue7 6
   ```

Enabling Buffer Statistics Tracking

This functionality is supported on the S6000 platform.

You can enable the tracking mechanism for statistical values of buffer spaces at a global level to be applicable throughout the system. By default, this capability to monitor buffer statistics is disabled. The buffer statistics tracking utility operates in the max use count mode that enables the collection of maximum values of counters.

To configure the buffer statistics tracking utility, perform the following:

Enable the buffer statistics tracking utility and enter the Buffer Statistics Snapshot configuration mode.

CONFIGURATION mode

```
FTOS(conf)#buffer-stats-snapshot
```

```
FTOS(conf)#no disable
```

You must enable this utility to be able to configure the parameters for buffer statistics tracking. By default, buffer statistics tracking is disabled.

Classifying Packets Based on a Combination of DSCP Code Points and VLAN IDs

This functionality is supported on the S4810, S4820T, Z9000, and MXL platforms.

You can configure a classifier map that contains both the DSCP and MAC VLAN IDs as parameters for performing a filtering of packets that are received before they are forwarded or dropped. You can now specify both DSCP—IP packet classification (Layer 3 headers) and Dot1p—(IEEE 802.1p) Packet classification (Layer 2 headers) as match criteria in a Layer 3 class map.

The type of the class map is determined during the creation of a class map. In releases of Dell Networking OS earlier than Release 9.2.0.0, you can configure only dot1p value as the filter criterion in Layer 2 class maps and DSCP value as the filter parameter in Layer 3 class maps. Classifying packets using both the Layer 2 attribute, dot1p value or MAC VLAN, in a Layer 2 class map and the Layer 3 attribute, DSCP value, in a Layer 3 class map is also possible. However, it was not possible to configure both dot1p or MAC VLAN, and DSCP values in the same L2 or L3 class map.

All class maps are Layer 3 by default. You can now configure a Layer 3 class map to differentiate traffic according to the IP VLAN value and the DSCP value. You can use the `match ip vlan` *`vlan-id`* command in Class Map Input Configuration mode to specify a match criterion for a class map based on a VLAN ID. You can attach this class map with a policy map, and associate the policy map with a service queue. When you link class-maps to queues using the `service-queue` command, Dell Networking OS matches the class-maps according to queue priority (queue numbers closer to 0 have lower priorities).

To create IP VLAN and DSCP values as match criteria in a Layer 3 class map, and to associate the class map with a policy map that is linked to a service queue, perform the following:

1. Create a match-any or a match-all Layer 3 class map, depending on whether you want the packets to meet all or any of the match criteria to be a member of the class. By default, a Layer 3 class map is created if you do not enter the layer2 option with the class-map command. When you create a class map, you enter the Class Map configuration mode.

   CONFIGURATION mode

   ```
   Dell (conf)#class-map match-all pp_classmap
   ```

2. Use a differentiated services code point (DSCP) value as a match criterion.

   CLASS-MAP mode

   ```
   Dell (conf-class-map)#match ipdscp 5
   ```

3. Configure a match criterion for a class map based on VLAN ID.

   CLASS-MAP mode

   ```
   Dell (conf-class-map)#match ip vlan 5
   ```

4. Create a QoS input policy on the device.

   CONFIGURATION mode

   ```
   Dell(conf)#qos-policy-input pp_qospolicy
   ```

5. Specify the DSCP value to be set on the matched traffic.

   QOS-POLICY-IN mode

   ```
   Dell(conf-qos-policy-in)#set ip-dscp 5
   ```

6. Create an input policy map.

   CONFIGURATION mode

   ```
   Dell(conf)#policy-map-input pp_policmap
   ```

7. Create a service queue to associate the class map and QoS policy map.

   POLICY-MAP mode

   ```
   Dell(conf-policy-map-in)#service-queue 0 class-map pp_classmap qos-policy
   pp_qospolicy
   ```

rate shape

Define the rate-shaping method to be either as a measure of bytes or packets for each of the hierarchical QoS (HQoS) nodes at the leaf level to be applied to each queue.

S6000

Syntax

```
[no] rate-shape [kbps | pps] peak-rate [burst-KB/Packets]
[committed [kbps | pps] committed-rate [burst-KB/Packets]]
```

Parameters

kbps	Enter the keyword kbps to specify the committed rate limit in Kilobits per second (Kbps). Specify this value as a multiple of 64. The range is from 0 to 40000000. The default granularity is Megabits per second (Mbps).
pps	Enter the keyword pps to specify the committed rate limit in packets per second (pps). The range is from 1 to 268000000. The default granularity is Megabits per second (Mbps).

committed-rate	Define the committed rate, which is the guaranteed or minimum confirmed rate for the packets. Specify this value as a multiple of 64 if you specify the committed rate in Kbps. The range is from 0 to 40000000 for Kbps. The range is from 1 to 268000000 for pps. The range is from 0 to 40000 for Mbps (which is the default measure for rate limits if you do not explicitly configure Kbps or pps)
burst-KB	(OPTIONAL) Enter the committed burst size in KB. The range is from 0 to 10000. The default is 50 KB. The default peak burst is regarded as the same value as the configured committed burst size.
Packets	(OPTIONAL) Enter the committed burst size as a count of packets. The range is from 1 to 1073000. The default is 50 packets. The default peak rate is regarded as the same value as the configured committed rate.
peak-rate	Define the peak rate, which is the guaranteed or minimum confirmed rate for the packets. Specify this value as a multiple of 64 if you specify the peak rate in Kbps. The range is from 0 to 40000000 for Kbps. The range is from 1 to 268000000 for pps. The range is from 0 to 40000 for Mbps (which is the default measure for rate limits if you do not explicitly configure Kbps or pps)
kbps	Enter the keyword kbps to specify the peak rate limit in Kilobits per second (Kbps). Specify this value as a multiple of 64. The range is from 0 to 40000000. The default granularity is Megabits per second (Mbps).
pps	Enter the keyword pps to specify the peak rate limit in packets per second (pps). The range is from 1 to 268000000. The default granularity is Megabits per second (Mbps).
peak-rate	Define the peak rate, which is the guaranteed or minimum confirmed rate for the packets. Specify this value as a multiple of 64 if you specify the peak rate in Kbps. The range is from 0 to 40000000 for Kbps. The range is from 1 to 268000000 for pps. The range is from 0 to 40000 for Mbps (which is the default measure for rate limits if you do not explicitly configure Kbps or pps)
burst-KB	(OPTIONAL) Enter the peak burst size in KB. The range is from 0 to 10000. The default is 50 KB.
Packets	(OPTIONAL) Enter the peak burst size as a count of packets. The range is from 1 to 1073000. The default is 50 packets. The default peak rate is regarded as the same value as the configured peak rate.

Default Granularity for rate is Mbps unless you use the kbps option.

Command Modes	CONFIGURATION
Command History	**Version 9.3.0.0** Added support for committed rate and committed burst size, and for configuration of rate limits on the S6000 platform.

Usage Information

If you specify the pps keyword after the rate-shape command, the peak rate, peak burst, committed rate and committed burst are all considered to be values as a measure of packets. If you do not specify the `pps` or `kbps` keyword, the peak and committed rate settings are considered to be values in Mbps. Similarly, if you enter the `kbps` keyword, the peak and committed rate settings are treated as values in Kbps.

You cannot configure the committed rate settings to use a different metric or unit from the metric that is set for peak rate attributes because when you use the `rate-shape kbps` command, it denotes the metric for peak and committed rate attributes). Similarly, if you use the `rate-shape pps` option , it denotes the metric for peak rate and committed rate attributes.

If you attempt to define the committed rate to be less than the peak rate, an error message is displayed stating that the peak rate cannot be lower than the committed rate. You can configure all the rate shaping parameters to be either in bytes or packets measure for each queue. The rate and burst parameters for both minimum and maximum settings for a queue can be either in packets or bytes. You cannot configure some of rate shaping attributes to be in bytes measure and the remaining rate shaping attributes to be in packets measure; all the rate shaping attributes must contain the same metric or unit of measure.

Example

```
Dell (conf-qos-policy-out) #rate-shape pps 100 100 peak pps
1000 200

Dell (conf-qos-policy-out) #rate-shape kbps 1024 100 peak kbps
102400 75

Dell (conf-qos-policy-out) # rate-shape 100 100 peak 1000 750

Dell(conf-qos-policy-in)#rate-police 100 25 peak 80 500
% Error: Peak rate cannot be less than committed rate.
```

buffer-stats-snapshot

Enable the buffer statistics tracking utility and enter the Buffer Statistics Snapshot configuration mode. You must enable this utility to be able to configure the parameters for buffer statistics tracking. This utility is supported on the S6000 platform.

S6000

Syntax

```
[No] buffer-stats-snapshot
```

To disable the buffer statistics tracking utility, enter the `disable` command from the BUFFER-STATS-SNAPSHOT mode.

Default	By default, buffer statistics tracking is disabled.
Command Modes	CONFIGURATION mode
Command History	**Version 9.3.0.0** Introduced on the S6000 platform
Usage Information	Only the software-based trigger for retrieving and calculating the snapshots of the statistical counters of the buffer space is supported. Collection of snapshots of buffer statistical counters based on hardware threshold settings is not supported, which can be used to specify the type of hardware threshold and the threshold profile templates.

Example

```
Dell(conf)#buffer-stats-snapshot
Dell(conf-buffer-stats-snapshot)#?
disable                 Disable buffer-stats-snapshot globally
end                     Exit from configuration mode
exit                    Exit from buffer-stats-snapshot
configuration mode
no                      Negate a command or set its defaults
show                    Show buffer-stats-snapshot
configuration
Dell(conf-buffer-stats-snapshot)#no disable
Dell(conf-buffer-stats-snapshot)#show configuration
!
buffer-stats-snapshot
 no disable
```

service-class buffer shared-threshold-weight

Create a service class and associate the threshold weight of the shared buffer with each of the queues per port in the egress direction. A global buffer pool that is a shared buffer pool accessed by multiple queues when the minimum guaranteed buffers for the queue are consumed can be configured on the S6000 and Z9000 platforms.

S6000Z9000

Syntax

```
[No] Service-class buffer shared-threshold-weight {[queue0
number] || [queue1 number] || [queue2 number] || [queue3
number] || [queue4 number] || [queue5 number] || [queue6
number] || [queue7 number]}
```

Parameters

service-class	Define the mapping between the service class and policy-based QoS or routing
buffer	Define the shared buffer settings
shared-threshold-weight	Specify the weight of a queue for the shared buffer space
queue 0 to queue 7	Specify the queue number to which the WRED parameters apply

number	Enter a weight for the queue on the shared buffer as a number in the range of 1 to 11.

Default The default threshold weight on the shared buffer for each queue is 9. Therefore, each queue can consume up to 66.67 percent of available shared buffer by default.

Command Modes INTERFACE mode

Command History

Version 9.3.0.0	Introduced on the S6000 platform

Usage Information You can configure all the data queues. For S6000, you can configure queues 0-7.

The following table describes the mapping between the threshold weight of the shared buffer on the queue and the percentage of available shared buffer that is used by the queue for each of the corresponding threshold weights of the shared buffer:

shared-threshold-weight on the queue	% of available shared buffer that can be consumed by the queue.
0	No dynamic sharing; shared buffer = 0.
1	0.77%
2	1.54%
3	3.03%
4	5.88%
5	11.11%
6	20%
7	33.33%
8	50%
9	66.67%
10	80%
11	88.89%

Example

```
Dell(conf-if-te-0/8)#Service-class buffer shared-threshold-
weight queue5 4 queue7 6
```

wred weight

Configure the weight factor for computation of average-queue size. This weight value applies to front-end ports. This mechanism to configure a weight for WRED and ECN functionality for front-end ports is supported on the S6000 and Z9000 platforms.

S6000Z9000

Syntax `[no] wred weight number`

Parameters

weight	Define the weight factor to be used for computation of the WRED average-queue size to either enable WRED to discard packets or cause ECN to mark packets that exceed the minimum threshold configured. This setting applies to front-end ports only.
number	Enter the weight as a number to be used to calculate the average-queue size. The range is 1 to 15. The default is 0.

Default The default weight is zero.

Command Modes QOS-POLICY-OUT mode

Command History

Version 9.3.0.0	Introduced on the S6000 and Z9000 platforms

Usage Information If the average queue size is more than the maximum threshold of WRED, the packet is dropped. If the average queue size is between the minimum and maximum threshold values, the decision to drop or queue the packet is taken based on the packet drop probability. The probability that a packet is dropped depends on the minimum threshold, maximum threshold, and mark probability denominator.

Example

```
FTOS (conf-qos-policy-out) # wred weight 5
```

service-class wred

The mechanism to configure a weight factor for WRED and ECN functionality for backplane ports is supported on the Z9000 platform. Also, this mechanism to configure a weight for WRED and ECN functionality for front-end ports is supported on the Z9000 platforms. Create a weighted random early detection (WRED) profile and ECN functionality per queue granularity for backplane ports, and attach the WRED profile with a service class. You can enable or disable these parameters for each queue and specify minimum and maximum buffer thresholds for each color-coding of the packets. Also, you can specify the maximum drop rate percentage for yellow and green profiles. The per-queue profile configured is applicable to all the backplane ports.

Z9000

Syntax

```
[No] service-class wred {green | weight | yellow} {[queue0
number/string] || [queue1 number/string] || [queue2 number/
```

string] || [queue3 *number/string*] || [queue4 *number/string*] || [queue5 *number/string*] || [queue6 *number/string*] || [queue7 *number/string*]}{backplane}

Parameters

service-class	Define the mapping between the service class and policy-based QoS or routing
wred	Specify WRED curve parameters for a queue
green	Specify green (low) drop precedence to a queue
weight	Specify a weight factor to a queue
yellow	Specify yellow (medium) drop precedence to a queue
queue 0 to queue 7	Specify the queue number to which the WRED parameters apply
number	Enter a weight for the queue as a number in the range of 1 to 15. This parameter applies only if you specify the green or yellow drop precedence.
string	Enter the WRED profile name. It is a string of up to 32 characters. Or use one of the five pre-defined WRED profile names. Pre-defined Profiles: wred_drop, wred-ge_y, wred_ge_g, wred_teng_y, wred_teng_. This parameter applies only if you specify a weight factor.
backplane	Specify that the WRED weight and profile configured for each queue apply to backplane ports

Default All queues on backplane ports operate in tail-drop (best-effort traffic) mode by default. There is no default WRED green or yellow profile. The default weight is 0.

Command Modes QOS-POLICY-OUT mode

Command History

Version 9.3.0.0	Introduced on the Z9000 platform

Usage Information You can configure all the data queues. For Z9000, you can configure queues 0-3. WRED profile contains a set of characteristics, such as the minimum and maximum WRED thresholds and the maximum drop rate. You can add and remove WRED parameters for one or more queues by using the command in a single line. All of the configured attributes apply to all the backplane ports and are for each queue. To assign drop precedence to green or yellow traffic, use this command. If there is no honoring enabled on the input, all the traffic defaults to green drop precedence.

Example

```
Dell(conf-wred) #wred thresh-1
Dell(conf-wred) #threshold min 100 max 200 max-drop-rate 40
Dell(conf-wred) #wred thresh-2
Dell(conf-wred) #threshold min 300 max 400 max-drop-rate 80
Dell(conf) #service-class wred green queue5 thresh-1 queue7
thresh-2 backplane
Dell(conf) #service-class wred yellow queue1 thresh-2 queue3
thresh-1 backplane
```

```
Dell(conf) #service-class wred weight queue0 11 queue6 4
queue7 9 backplane
```

service-pool wred

A global buffer pool that is a shared buffer pool accessed by multiple queues when the minimum guaranteed buffers for the queue are consumed can be configured on the S6000 and Z9000 platforms.

Create a global buffer pool that is a shared buffer pool accessed by multiple queues when the minimum guaranteed buffers for the queue are consumed. S4810, S4820T, S6000, and Z9000 platforms support four global service-pools in the egress direction. Two service pools are used—one for lossy queues and the other for lossless (priority-based flow control (PFC)) queues. You can enable WRED and ECN configuration on the global service-pools. You can define WRED profiles and weight on each of the global service-pools for both lossy and lossless (PFC) service-pools.

S6000Z9000

Syntax

`[No] buffer-pool wred {green | weight | yellow} {[pool0 number/string] || [pool1 number/string]}`

Parameters

buffer-pool	Define the mapping between the service class and policy-based QoS or routing
wred	Specify WRED curve parameters for a queue
green	Specify green (low) drop precedence to a queue
weight	Specify a weight factor to a queue
yellow	Specify yellow (medium) drop precedence to a queue
pool0	Service-pool buffer 1 (default service-pool for PFC traffic)
pool1	Service-pool buffer 0 (default service-pool for lossy traffic)
number	Enter a weight for the queue as a number in the range of 1 to 15. This parameter applies only if you specify the green or yellow drop precedence.
string	Enter the WRED profile name. It is a string of up to 32 characters. Or use one of the five pre-defined WRED profile names. Pre-defined Profiles: wred_drop, wred-ge_y, wred_ge_g, wred_teng_y, wred_teng_. This parameter applies only if you specify a weight factor.

Default

All queues on backplane ports operate in tail-drop (best-effort traffic) mode by default. There is no default WRED green or yellow profile. The default weight is 0.

Command Modes

CONFIGURATION mode

Command History

Version 9.3.0.0	Introduced on the S6000 and Z9000 platforms

Usage Information

You can configure only service pools 0 and 1 because the Dell Networking OS uses only these two service pools. The pool, service0, is used for lossy queues and the pool, service1, is used for lossless (PFC) queues in all the platforms.

You can configure the weight for the WRED average queue size for service1 on the S6000 Switch, which is the only platform in which PFC is supported for this service pool. On the Z9000 Switch, only service0 can be configured because it does not support PFC.

A WRED profile contains a set of attributes, such as the minimum and maximum threshold values, and the maximum drop rate for the received packets. You can add or remove WRED parameter configurations for one or more shared service pools using a single command. The `buffer-pool wred` command is similar in usage and working to the `service-class bandwidth-percentage` *queue-id* command.

Example

```
Dell(conf-wred) #wred thresh-1
Dell(conf-wred) #threshold min 100 max 200 max-drop-rate 40

Dell(conf-wred) #wred thresh-2
Dell(conf-wred) #threshold min 300 max 400 max-drop-rate 80

Dell(conf) #service-pool wred green pool0 thresh-1 pool1
thresh-2
Dell(conf) #service-pool wred yellow pool0 thresh-3 pool1
thresh-4
Dell(conf) #service-pool wred weight pool0 11 pool1 4
```

service-class wred

Create a service class and assign ECN marking for different queues on backplane ports to the service class. This functionality can be configured on the Z9000 platforms.

Z9000

Syntax

```
[No] service-class wred ecn queues-list {backplane}
```

Parameters

service-class	Define the mapping between the service class and policy-based QoS or routing
wred	Associate WRED with ECN to mark packets instead of dropping them
ecn	Cause explicit congestion notification (ECN) to be used to indicate network congestion, rather than dropping packets, queues-list Enter the queue numbers, either as individual queue numbers separated by commas or as an inclusive list separating the starting and ending queue numbers with a hyphen
backplane	Specify that the ECN marking configured for each queue applies to backplane ports

Default

By default, ECN marking is disabled on all queues.

Command Modes	CONFIGURATION mode
Command History	**Version 9.3.0.0** Introduced on the S6000 and Z9000 platforms
Usage Information	You can add or remove ECN marking configuration on a list of queues on all backplane ports. All of the configured attributes apply to all the backplane ports and are for each queue. You can configure all the data queues. For Z9000, you can configure queues 0-3. By default, ECN marking is disabled on all queues. When you enable wred-ecn, and the number of packets in the queue is below the minimum threshold, packets are transmitted per the usual WRED treatment. When you enable wred-ecn, and the number of packets in the queue is between the minimum threshold and the maximum threshold, one of the following three scenarios can occur: • If the transmission endpoints are ECN-capable and traffic is congested, and the WRED algorithm determines that the packet should have been dropped based on the drop probability, the packet is transmitted and marked so the routers know the system is congested and can slow transmission rates. • If neither endpoint is ECN-capable, the packet may be dropped based on the WRED drop probability. This behavior is the identical treatment that a packet receives when WRED is enabled without ECN configured on the router. • If the network is experiencing congestion, the packet is transmitted. No further marking is required. When you enable wred-ecn, and the number of packets in the queue is above the maximum threshold, packets are dropped based on the drop probability. This behavior is the identical treatment a packet receives when WRED is enabled without ECN configured on the router.
Example	`Dellconf) #service-class wred ecn 0, 3-5, 7 backplane`

service-class wred ecn

Create a service class and assign ECN marking for different queues on backplane ports to the service class. This functionality can be configured on the Z9000 platforms.

Z9000

Syntax `[No] service-class wred ecn queues-list {backplane}`

Parameters

service-class	Define the mapping between the service class and policy-based QoS or routing
wred	Associate WRED with ECN to mark packets instead of dropping them
ecn	Cause explicit congestion notification (ECN) to be used to indicate network congestion, rather than dropping packets, queues-list Enter the queue numbers, either as individual queue numbers separated by commas or as an inclusive list

		separating the starting and ending queue numbers with a hyphen
	backplane	Specify that the ECN marking configured for each queue applies to backplane ports

Default By default, ECN marking is disabled on all queues.

Command Modes CONFIGURATION mode

Command History

Version 9.3.0.0	Introduced on the Z9000 platform

Usage Information You can add or remove ECN marking configuration on a list of queues on all backplane ports. All of the configured attributes apply to all the backplane ports and are for each queue. You can configure all the data queues. For Z9000, you can configure queues 0-3. For S6000, you can configure queues 0-7. By default, ECN marking is disabled on all queues. When you enable wred-ecn, and the number of packets in the queue is below the minimum threshold, packets are transmitted per the usual WRED treatment. When you enable wred-ecn, and the number of packets in the queue is between the minimum threshold and the maximum threshold, one of the following three scenarios can occur:

- If the transmission endpoints are ECN-capable and traffic is congested, and the WRED algorithm determines that the packet should have been dropped based on the drop probability, the packet is transmitted and marked so the routers know the system is congested and can slow transmission rates.
- If neither endpoint is ECN-capable, the packet may be dropped based on the WRED drop probability. This behavior is the identical treatment that a packet receives when WRED is enabled without ECN configured on the router.
- If the network is experiencing congestion, the packet is transmitted. No further marking is required. When you enable wred-ecn, and the number of packets in the queue is above the maximum threshold, packets are dropped based on the drop probability. This behavior is the identical treatment a packet receives when WRED is enabled without ECN configured on the router.

Example

```
Dellconf) #service-class wred ecn 0, 3-5, 7 backplane
```

show hardware stack-unit buffer

Display the counters for the specified port, minimum guaranteed buffer of a priority-group, and the shared buffer. This command is supported on the S6000 platform.

Syntax

```
show hardware stack-unit <0-5> buffer unit <0-0> port {1-104 |
all} priority-group <0 -7> buffer-info
```

Parameters

stack-unit *<0–5>*	Enter the keywords `stack-unit` to select a particular stack member and then enter one of the following command options to display a collection of data based on the option entered. The range is 0 to 5.

buffer unit <*0–5*>	Enter the keyword `buffer`. To display the total buffer statistics for the stack unit, enter the keyword total-buffer. To display the buffer statistics for a specific unit, enter the keyword *unit* and a unit number 0 .
port	To display the buffer statistics for a specific port, enter the keyword `port` and a port number from 1 to 64.
all	Display buffer statistics for all ports
priority-group	Identifier of the priority group in the range of 0 to 7.
buffer-info	To display total buffer information for the port, enter the keywords `buffer-info`. To display a queue range, enter 0 to 14 for a specfic queue or all.

Command Modes

EXEC

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S6000 platform.

Usage Information

The following table describes the fields in the output of the `show` command:

Field	Description
Buffer Accounting Stats for Unit 0 Port 1 priority-group 0	Displays the counters that are calculated by the buffer statistics tracking method for each port per priority group on a particular stack member.
Max Shared Limit	Maximum shared buffer space allotted to the specific port for the corresponding stack unit
Default Packet Buffer allocate for the priority-group	The default packet buffer size in KB that is associated with the particular priority group
Accounted Packet Buffer	Shared buffer space that is in use by the packets

Example

```
FTOS# show hardware stack-unit 0 buffer unit 0 port 1 priority-
group 0 buffer-info

-----  Buffer Accounting Stats for Unit 0 Port 1 priority-
group 0 -----
 Maximum Shared Limit: 0
 Default Packet Buffer allocate for the priority-group: 61440
 Accounted Packet Buffer: 0
```

show hardware stack-unit buffer-stats-snapshot

View the buffer statistics tracking resource information without polling details and historical snapshots. This command is supported on the S6000 platform.

Syntax

```
show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource X
```

Parameters

stack-unit *number*	Unique ID of the stack unit to select a particular stack member and then enter one of the following command options to display a collection of data based on the option entered. The range is 0 to 11.
buffer-stats-snapshot unit *number*	Display the historical snapshot of buffer statistical values unit Enter the keyword `unit` along with a port-pipe number, then the keyword counters to display the counters on the selected port-pipe. The range is 0 to 0.
port	resource X Buffer and traffic manager resources usage, where X can be one of the following: • All - Ingress and Egress resources snapshots • Port {id |all} queue {all} - egress queue-level snapshot for both unicast and multicast packets • Port {id |all} queue ucast {id | all} - egress queue-level snapshot for unicast packets only • Port {id |all} queue mcast {id | all} - egress queue-level snapshot for multicast packets only • Port {id |all} prio-group {id | all} - ingress priority-group level snapshot

Command Modes

EXEC

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S6000 platform.

Usage Information

The following information is displayed depending on whether the historical snapshot of buffer statistics is needed for all ports, per-port per-queue, or a priority group.

- All – Displays all resources on ingress and egress for each of the port, queue.
- Port-Queue ucast/mcast – Displays the total unicast/multicast buffer usage on per-port per-queue basis. For CPU port, counters for queues 0 – 11 are displayed and no differentiation is made between unicast and multicast queues.
- Port- Priority-Group – Displays the shared space counters usage, head-room space counters per ingress port on per-priority- group granularity.

When the buffer-stats-snapshot is disabled, an informational message is dispayed to this effect when you attempt to view the buffer statistics tracking resource information without polling details and historical snapshots.

Example

```
FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 1 queue ucast all
Stack-unit: 0 unit: 0 port: 1 (interface Fo 0/0)
--------------------------------------
Q# TYPE     Q#     TOTAL BUFFERED CELLS
--------------------------------------
UCAST       0      0
UCAST       1      0
UCAST       2      0
UCAST       3      0
UCAST       4      0
UCAST       5      0
UCAST       6      0
UCAST       7      0
UCAST       8      0
UCAST       9      0
UCAST       10     0
UCAST       11     0


FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 5 queue all

Stack-unit 0 unit 0 port 5 (interface te 0/4)

------------------------------------------------------------------
-----------------------------------------------
Q# TYPE     Q#      TOTAL BUFFERED CELLS
------------------------------------------------------------------
-----------------------------------------------
UCAST        5                   4
UCAST        6                   8
UCAST        11                 1
MCAST       4                    11
```

Only the queues for which the buffer cell consumption is not zero are displayed. If an egress buffer is not present on any of the queues on port 5, the following sample output is displayed:

```
FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 5 queue all

Stack-unit 0 unit 0 port 5 (interface te 0/4)


------------------------------------------------------------------
-----------------------------------------------
Q# TYPE     Q#      TOTAL BUFFERED CELLS
------------------------------------------------------------------
-----------------------------------------------

FTOS#show hardware stack-unit 0 buffer-stats-snapshot unit 0
resource port 5 prio-group all

Stack-unit 0 unit 0 port 5 (interface te 0/4)
```

```
-----------------------------------------------------------------
-----------------------------------------------
PG#     SHARED CELLS      HEADROOM CELLS
-----------------------------------------------------------------
-----------------------------------------------
6           1000
5
7           3                                   0
```

show hardware stack-unit buffer-stats-snapshot (Total Buffer Information)

View the buffer statistics tracking resource information depending on the type of buffer information, such as device-level details, port-level counters, queue-based snapshots, or priority group-level snapshot in the egress and ingress direction of traffic. This command is supported on the S6000 platform.

Syntax

```
show hardware stack-unit <0–11> buffer-stats-snapshot unit <0–0> buffer-info x
```

Parameters

stack-unit <*0–11*>
Unique ID of the stack unit to select a particular stack member and then enter one of the following command options to display a collection of data based on the option entered. The range is 0 to 11.

buffer-stats-snapshot unit *number*
Display the historical snapshot of buffer statistical values unit Enter the keyword `unit` along with a port-pipe number, then the keyword counters to display the counters on the selected port-pipe. The range is 0 to 0.

buffer-info
buffer-info Displays total buffer information for a group, where *x* can be one of the following:

- All - Displays ingress and egress device, port, and queue snapshots
- Port {id |all} Displays both ingress and egress port-level snapshot
- Port ingress {id |all} Displays ingress port-level snapshot
- Port egress {id |all) Displays egress port-level snapshot
- Port {id |all} queue {all} - egress queue-level snapshot for both unicast and multicast packets
- Port {id |all} queue ucast {id | all} - egress queue-level snapshot for unicast packets only
- Port {id |all} queue mcast {id | all} - egress queue-level snapshot for multicast packets only
- Port {id |all} prio-group {id | all} - ingress priority-group level snapshot

Command Modes	EXEC EXEC Privilege
Command History	**Version 9.3.0.0** Introduced on the S6000 platform.
Usage Information	The following information is displayed based on the buffer-info type, such as device-level details, port-level counters, queue-based snapshots, or priority group-level snapshot in the egress and ingress direction of traffic: • Device-ingress – Displays total buffer accounting usage for the unit. • Device-egress –Display total buffer usage for the unit, total multicast buffer usage for the unit and also on per-service-pool basis. Counters will be displayed for the 2 service-pools – one for normal traffic and other for DCB traffic. • Port-ingress – Displays the total buffer accounting usage for the ingress port. • Port-egress – Displays the total unicast buffer usage, total multicast buffer usage separately for the egress port. • Port-Queue ucast/mcast – Displays the total unicast/multicast buffer usage on per-port per-queue basis. For CPU port, counters for queues 0 – 11 will be displayed and there is no differentiation between unicast and multicast queues. • Port- Priority-Group – Displays the shared space usage counters usage, head-room space counters per ingress port on per-priority-group granularity. When the buffer-stats-snapshot is disabled, the folloing informational message is displayed when you run the show command: `%Info: Buffer-stats-snapshot feature is disabled.`

15

Management Port Media Converter

The capability to configure management media port converter is supported on the S6000 platform.

Starting with Dell Networking OS Release 9.3.0.0, copper Ethernet network management connectivity for power distribution units (PDUs) is supported, without the need to provision an additional network switch. A unique, dedicated special media converter is provided and it can be inserted in to any front-panel 40G interface. This converter supports 10M, 100M, 1G Ethernet speeds. Although these ports are used for management connectivity, the OS considers the traffic traversing through these ports as traffic passing through any other data port.

To support the connectivity between the QSFP port of S6000 and the Base-T port of the PDU, a special breakout Media Converter has been designed, with one end as the QSFP module and the other end with 4 RJ45 connectors to provide 4x 10/100/1000Base-T functionality.

When the media converter is inserted into the device and is detected by the switch, you can perform the following tasks:

Configure the 10, 100 or 1000Base-T operation through an Serial Gigabit Media Independent Interface (SGMII) interface.

Monitor the link

View the statistical counters

The S6000 platform does not support SGMII mode with autonegotiation. To negotiate with the peer PDU, a new physical layer (PHY) component is provided. The Media converter supports SGMII autonegotiation with peer PDUs.

Also, the S6000 platform does not support half-duplex mode at any speed and only full-duplex mode of operation is supported. Therefore, support for full-duplex is made available for speeds of 10M and 100M, although PDUs support half-duplex mechanism.

Management Port Media Converter Components

The media converter is a 40G to 4 x 10/100/1000Base-T converter. The following are the elements of the media converter:

- Electrically erasable programmable read-only memory (EEPROM) in Dell standard to provide vendor details: I2C Address : 0x50 (7-bit)
- 10/100/1000 Base-T Gigabit Ethernet transceiver-4 per QSFP; 1 chip per fanout port ; I2C Address: 0x52 (7-bit)
- I2C Multiplexer- PCA9546A: 4 channel I2C bus switch. 1 per QSFP.
  - I2C Address: 0x74 (7-bit)
  - I2C Mux is used to select one of the four PHY54616s.

All three components of EEPROM, BCMPHY, I2C Mux are accessed using I2C, with their unique addresses.

Working of the Management Port Media Converter

QSFP EEPROM content is accessed to detect this special media converter and based on this identification performed, a new optic type is used to differentiate this optic type from the other optics. This new optic type is displayed in the appropriate `show` commands and for enabling other functionalities. EEPROM contents are displayed in the output of the 'show interface transceiver' command. The optic functions in the same manner as any other fanout QSFP 4x10G cable, with the exception that this optic needs the speed to be manually configured in the software.

Because of the optic type that is present in the media converter to connect the QSFP port with the PDU requiring the speed to be configured, you can use the speed command to configure each of the 4 fanout ports with the required speed of 1G, 100M, or 10M. Because the S6000 platform does not provide the autonegotiation in SGMII mode, it is required that you configure the speed of the individual 4 fanout ports based on the peer PDU's configuration by using the speed command. By default, each port runs at a speed of 1G during the time of insertion of the optic.

The speed command is available for configuration only when the management optic is inserted in to a quad-mode interface. Because this optic cannot be used in a non-fanout port, you must ensure that fanout is enabled to use this optic, similar to how it is enabled for other 4x10G fanout cables. All existing commands and configuration settings for that port are retained with the additional speed command setting.

Based on the optic detection and fanout mode configuration, you can use the speed command. This command is activated for configuration on optic insertion and removed on optic removal. In non-quad mode, the speed CLI is not enabled (if you did not configure the stack-unit stack-unit port number portmode quad command). The media converter `speed` CLI is enabled only on fanout configured ports. If an optic is inserted in to a non-fanout port, there is no change in the functionality and you must configure fanout and reload the device to use this optic. When the management (MGMT) optic is inserted, the default speed is configured and the speed CLI is enabled. When MGMT optic is removed, the speed CLI is disabled and its related settings are removed from the running configuration.

The following is the portion of the sample output of the show config command when the management optic is inserted:

```
FTOS(conf-if-te-2/10)#show config
interface TenGigabitEthernet 2/10
 speed 100
 no shutdown
```

The following is the sample output of the show interface command when the management optic is inserted:

```
FTOS>show interface tengig 0/0
TenGigabitEthernet 0/0 is up, line protocol is up
Hardware is DellForce10Eth, address is 90:b1:1c:f4:9a:fa
    Current address is 90:b1:1c:f4:9a:fa
Pluggable media present, QSFP type is 40GBASE-XXXX
LineSpeed 100 Mbit
```

The following system log messages are displayed for optic insertion and removal, which are identical to the messages displayed for other QSFP optics:

On Optic Insertion:

```
00:00:40: %S6000:0 %IFAGT-5-INSERT_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 0
00:00:40: %S6000:0 %IFAGT-5-INSERT_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 1
00:00:40: %S6000:0 %IFAGT-5-INSERT_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 2
```

```
00:00:40: %S6000:0 %IFAGT-5-INSERT_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 3
```

On Optic Removal:

```
00:04:41: %S6000:0 %IFAGT-5-REMOVE_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 0
00:04:41: %S6000:0 %IFAGT-5-REMOVE_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 1
00:04:41: %S6000:0 %IFAGT-5-REMOVE_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 2
00:04:41: %S6000:0 %IFAGT-5-REMOVE_OPTICS_QSFP: Optics QSFP inserted in slot 0
port 3
```

The following is the sample output of the show interface media command when the management optic is inserted:

```
Slot   Port     Type        Media               Serial Number
F10Qualified
------------------------------------------------------------------------------
   0      0     QSFP        40GBASE-XXX         QB520116                 Yes
   0      1     QSFP        40GBASE-XXX         QB520116                 Yes
   0      2     QSFP        40GBASE-XXX         QB520116                 Yes
   0      3     QSFP        40GBASE-XXX         QB520116                 Yes
```

The following is the sample output of the show interface transceiver command when the management optic is inserted:

```
QSFP 0 Serial ID Base Fields
QSFP 0 Id                       = 0x0d
QSFP 0 Ext Id                   = 0x00
QSFP 0 Connector                = 0x0c
QSFP 0 Transceiver Code         = 0x04 0x00 0x00 0x00 0x00 0x00 0x00 0x00
QSFP 0 Encoding                 = 0x05
QSFP 0 Length(SFM)     Km       = 0x00
QSFP 0 Length(OM3)     2m       = 0x32
QSFP 0 Length(OM2)     1m       = 0x00
QSFP 0 Length(OM1)     1m       = 0x00
QSFP 0 Length(Copper)  1m       = 0x00
QSFP 0 Vendor Rev               = 0
```

Based on detecting the management optic type, the default settings are set in both the MAC and PHY elements.

On the MAC:

1. Set speed as 1G.
2. Duplex as Full
3. Interface type as SGMII
4. Autoneg off.

On the PHY:

1. Auto neg ON.
2. Full-duplex and half-duplex on.
3. Advertise 10M,100M,1G
4. Out of reset.

Online Insertion and Removal (OIR) of the Management Optic

The following table illustrates the system operations during various system states when the management optic is inserted and removed. It also describes the functioning of the system before the introduction of the management optic and the current working with the management optic type supported.

Table 6. Online Insertion and Removal (OIR) of the Management Optic

1.a	Run-time	No Optic present.	Insert MGMT optic	Speed CLI will be enabled. and default speed config applied.
S.No	During	Earlier Operation	Current Operation	System Behavior
1.b		Non – MGMT optic	Insert MGMT optic	As in 1.a
1.c		MGMT optic Present	Remove MGMT optic	Disable speed CLI and remove the speed config on that interface.
1.d		MGMT optic Removed	Reinsert MGMT optic	As in 1.a
1.e		MGMT optic Removed	Insert Non-MGMT optic	No change.
2.a	Bootup	MGMT optic removed during reload. Speed configs still present in startup-config.	Insert Non-MGMT optic	First time when Non-MGMT optic is inserted remove speed CLI configs for that port.
2.b		MGMT optic removed during reload. Speed configs still present in startup-config.	Insert MGMT optic	Enable speed CLI and update the earlier configured speed. S/W : Being 'speed' CLI is enabled only during Optic poller event, startup config might get applied too early. This has to be taken care.
2.c		10G before reload	40G after reload	Speed config will be removed, when the interface Te interface is deleted.

The following table describes the different scenarios of operations with the management optic type of the media port converter on the S6000 platform connected to the S50 device.

Table 7. Scenarios of Working of Management Optic Type

CASE NO	Speed	S6k T2				Bcm 54616s				S50				Remarks
		AN	SP	FD	HD	AN	SP	FD	HD	AN	SP	FD	HD	
Default config on both sides:														
1	1G	off	1G	ON	-	ON	Default Advertise: 10M, 100M, 1G	Default	Default	ON	Default Advertise: 10M, 100M, 1G	def	NA	Works only if auto ON at both sides

CASE NO	Speed	S6k T2				Bcm 54616s				S50				Remarks
		AN	SP	FD	HD	AN	SP	FD	HD	AN	SP	FD	HD	
Speed/Duplex change in peer side:														
Autoneg ON both sides, with config change in peer side														
2.a	100M FD	Off	100M	ON	-	On	Default	Default	Default	on	100M	Y/ Default	Y/ Default	Works
2.b	100M HD	Off	100M	ON	-	on	Default	Default	Default	on	100M	Off	Y	Works
2.c	10M FD	Off	10M	ON	-	On	Default	Default	Default	on	10M	Y/ Default	Y/ Default	Works
2.d	10M HD	Off	10M	ON	-	on	Default	Default	Default	on	10M	Off	Y	Works
Speed/Duplex change in s6k side:														
Autoneg ON both sides, with config change in mgmt. phy side														
3.a	100M FD	Off	100M	ON	-	On	100M	Y	Off	On	Default	Default	Default	Works
3.b	100M HD	Off	100M	ON	-	On	100M	off	Y	On	Default	Default	Default	Works
3.c	10M FD	Off	10M	ON	-	On	10M	Y	Off	On	Default	Default	Default	Works
3.d	10M HD	Off	10M	ON	-	On	10M	off	Y	On	Default	Default	Default	Works
Autoneg OFF at peer side, ON at mgmt. phy side.														
4.a	100M FD	Off	100M	ON	-	On	100M	Y	Off	off	100M	Y	Off	Links up in 100M HD in an=on side, as AN fail
4.b	100M HD	Off	100M	ON	-	On	100M	off	Y	off	100M	Off	Y	Works
4.c	10M FD	Off	10M	ON	-	On	10M	Y	Off	off	10M	Y	Off	Links up in 100M HD in an=on side, as AN fail
4.d	10M HD	Off	10M	ON	-	On	10M	off	Y	off	10M	Off	Y	Works
Autoneg OFF at mgmt. phy side , ON at peer side														
5.a	100M FD	Off	100M	ON	-	Off	100M	Y	Off	ON	100M	Y	Off	Links up in 100M HD in an=on side, as AN fail
5.b	100M HD	Off	100M	ON	-	Off	100M	off	Y	ON	100M	Off	Y	
5.c	10M FD	Off	10M	ON	-	Off	10M	Y	Off	ON	10M	Y	Off	Links up in 100M HD in

CASE NO	Speed	S6k T2				Bcm 54616s				S50				Remarks
		AN	SP	FD	HD	AN	SP	FD	HD	AN	SP	FD	HD	
														an=on side, as AN fail
5.d	10M HD	Off	10M	ON	-	Off	10M	off	Y	ON	10M	Off	Y	
Autoneg off on both sides, with manual config														
6.a	100M FD	Off	100M	ON	-	Off	100M	Y	Off	off	100M	Y	Off	Works
6.b	100M HD	Off	100M	ON	-	Off	100M	off	Y	off	100M	Off	Y	Works
6.c	10M FD	Off	10M	ON	-	Off	10M	Y	Off	off	10M	Y	Off	Works
6.d	10M HD	Off	10M	ON	-	Off	10M	off	Y	off	10M	Off	Y	works

16

Security for M I/O Aggregator

Security features are supported on the M I/O Aggregator.

This chapter describes several ways to provide access security to the Dell Networking system.

For details about all the commands described in this chapter, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

aaa authentication enable

Configure AAA Authentication method lists for user access to EXEC privilege mode (the "Enable" access).

Syntax

`aaa authentication enable {default | method-list-name} method [... method2]`

To return to the default setting, use the `no aaa authentication enable {default | method-list-name} method [... method2]` command.

Parameters

default	Enter the keyword `default` then the authentication methods to use as the default sequence of methods for the Enable login. The default is `default enable`.
method-list-name	Enter a text string (up to 16 characters long) to name the list of enabled authentication methods activated at login.
method	Enter one of the following methods: • `enable`: use the password the `enable password` command defines in CONFIGURATION mode. • `line`: use the password the `password` command defines in LINE mode. • `none`: no authentication. • `radius`: use the RADIUS servers configured with the `radius-server host` command. • `tacacs+`: use the TACACS+ server(s) configured with the `tacacs-server host` command.
... method2	(OPTIONAL) In the event of a "no response" from the first method, FTOS applies the next configured method.

Defaults Use the `enable` password.

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Usage Information

By default, the `Enable password` is used. If you configure `aaa authentication enable default`, FTOS uses the methods defined for `Enable access` instead.

Methods configured with the `aaa authentication enable` command are evaluated in the order they are configured. If authentication fails using the primary method, FTOS employs the second method (or third method, if necessary) automatically. For example, if the TACACS+ server is reachable, but the server key is invalid, FTOS proceeds to the next authentication method. The TACACS+ is incorrect, but the user is still authenticated by the secondary method.

Related Commands

enable password — changes the password for the `enable` command.

login authentication — enables AAA login authentication on the terminal lines.

password — creates a password.

radius-server host — specifies a RADIUS server host.

tacacs-server host — specifies a TACACS+ server host.

aaa authentication login

Configure AAA Authentication method lists for user access to EXEC mode (`Enable log-in`).

Syntax

`aaa authentication login {method-list-name | default} method [... method4]`

To return to the default setting, use the `no aaa authentication login {method-list-name | default}` command.

Parameters

method-list-name	Enter a text string (up to 16 characters long) as the name of a user-configured method list that can be applied to different lines.
default	Enter the keyword `default` to specify that the method list specified is the default method for all terminal lines.
method	Enter one of the following methods: • `enable`: use the password the `enable password` command defines in CONFIGURATION mode. • `line`: use the password the `password` command defines in LINE mode. • `none`: no authentication. • `radius`: use the RADIUS servers configured with the `radius-server host` command. • `tacacs+`: use the TACACS+ servers configured with the `tacacs-server host` command.
... method4	(OPTIONAL) Enter up to four additional methods. In the event of a "no response" from the first method, FTOS applies the next configured method (up to four configured methods).

Defaults	Not configured (that is, no authentication is performed).
Command Modes	CONFIGURATION
Command History	**Version 9.3.0.0** Introduced on the M I/O Aggregator.
Usage Information	By default, the locally configured username password is used. If you configure `aaa authentication login default`, FTOS uses the methods this command defines for login instead. Methods configured with the `aaa authentication login` command are evaluated in the order they are configured. If users encounter an error with the first method listed, FTOS applies the next method configured. If users fail the first method listed, no other methods are applied. The only exception is the local method. If the user's name is not listed in the local database, the next method is applied. If the correct user name/password combination is not entered, the user is not allowed access to the switch. **NOTE:** If authentication fails using the primary method, FTOS employs the second method (or third method, if necessary) automatically. For example, if the TACACS+ server is reachable, but the server key is invalid, FTOS proceeds to the next authentication method. The TACACS+ is incorrect, but the user is still authenticated by the secondary method. After configuring the `aaa authentication login` command, configure the `login authentication` command to enable the authentication scheme on terminal lines. Connections to the SSH server work with the following login mechanisms: local, radius, and tacacs.
Related Commands	login authentication — enables AAA login authentication on the terminal lines. password — creates a password. radius-server host — specifies a RADIUS server host. tacacs-server host — specifies a TACACS+ server host.

access-class

Restrict incoming connections to a particular IP address in a defined IP access control list (ACL).

Syntax	`access-class` *`access-list-name`* To delete a setting, use the `no access-class` command.
Parameters	***access-list-name*** Enter the name of an established IP Standard ACL.
Defaults	Not configured.
Command Modes	LINE

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

Related Commands

line — applies an authentication method list to the designated terminal lines.

ip access-list standard — names (or selects) a standard access list to filter based on the IP address.

ip access-list extended — names (or selects) an extended access list based on the IP addresses or protocols.

Authorization and Privilege Commands

To set command line authorization and privilege levels, use the following commands.

banner exec

Configure a message that is displayed when your enter EXEC mode.

Syntax

`banner exec c line c`

To delete a banner, use the `no banner exec` command.

Parameters

c	Enter the keywords `banner exec`, then enter a character delineator, represented here by the letter *c*. Press **ENTER**.
line	Enter a text string for your banner message ending the message with your delineator. In the following example, the delineator is a percent character (%); the banner message is "testing, testing".

Defaults No banner is displayed.

Command Modes CONFIGURATION

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

Usage Information After entering the banner login command, type one or more spaces and a delineator character. Enter the banner text then the second delineator character. When the user is connected to the router, if a message of the day banner is configured, it displays first. If no message of the day banner is configured, the login banner and prompt appear. After the user has logged in, the EXEC banner (if configured) displays.

Example

```
FTOS(conf)#banner exec ?
LINE c banner-text c, where 'c' is a delimiting character
FTOS(conf)#banner exec %
Enter TEXT message. End with the character '%'.
This is the banner%
FTOS(conf)#end
FTOS#exit
4d21h5m: %RPM0-P:CP %SEC-5-LOGOUT: Exec session is terminated
for user on line
console
```

```
This is the banner

FTOS con0 now available

Press RETURN to get started.
4d21h6m: %RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for
user on line
console
This is the banner
FTOS>
```

Related Commands

banner login — sets a banner for login connections to the system.

banner motd — sets a Message of the Day banner.

exec-banner — Enables the display of a text string when you enter EXEC mode.

line — enables and configures the console and virtual terminal lines to the system.

banner login

Set a banner to display when logging on to the system.

Syntax

```
banner login {keyboard-interactive | no keyboard-interactive}
[c line c]
```

Parameters

keyboard-interactive	Enter the keyword `keyboard-interactive` to require a carriage return (CR) to get the message banner prompt.
c	Enter a delineator character to specify the limits of the text banner. The delineator is a percent character (%).
line	Enter a text string for your text banner message ending the message with your delineator. The delineator is a percent character (%). Range: maximum of 50 lines, up to 255 characters per line

Defaults

No banner is configured and the CR is required when creating a banner.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Usage Information

After entering the banner login command, type one or more spaces and a delineator character. Enter the banner text then the second delineator character. When the user is connected to the router, if a message of the day banner is configured, it displays first. If no message of the day banner is configured, the login banner and prompt appear. After the user has logged in, the EXEC banner (if configured) displays.

Example

```
FTOS(conf)#banner login ?
keyboard-interactive Press enter key to get prompt
LINE c banner-text c, where 'c' is a delimiting character
FTOS(conf)#no banner login ?
keyboard-interactive Prompt will be displayed by default
```

```
<cr>
FTOS(conf)#banner login keyboard-interactive

Enter TEXT message. End with the character '%'.
This is the banner%
FTOS(conf)#end
FTOS#exit

13d21h9m: %RPM0-P:CP %SEC-5-LOGOUT: Exec session is terminated
for user on line console

This is the banner

FTOS con0 now available

Press RETURN to get started.
13d21h10m: %RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful
for user on line console
This is the banner
FTOS>
```

Related Commands

banner motd — sets a Message of the Day banner.

exec-banner — enables the display of a text string when you enter EXEC mode.

banner motd

Set a message of the day (MOTD) banner.

Syntax

```
banner motd c line c
```

Parameters

c	Enter a delineator character to specify the limits of the text banner. The delineator is a percent character (%).
line	Enter a text string for your MOTD banner the message with your delineator. The delineator is a percent character (%).

Defaults

No banner is configured.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Usage Information

After entering the banner login command, type one or more spaces and a delineator character. Enter the banner text then the second delineator character. When the user is connected to the router, if a message of the day banner is configured, it displays first. If no message of the day banner is configured, the login banner and prompt appear. After the user has logged in, the EXEC banner (if configured) displays.

Related Commands

banner exec — enables the display of a text string when you enter EXEC mode.

banner login — sets a banner to display after successful login to the system.

debug radius

View RADIUS transactions to assist with troubleshooting.

Syntax	`debug radius` To disable debugging of RADIUS, use the `no debug radius` command.	
Defaults	Disabled.	
Command Modes	EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

debug tacacs+

To assist with troubleshooting, view TACACS+ transactions.

Syntax	`debug tacacs+` To disable debugging of TACACS+, use the `no debug tacacs+` command.	
Defaults	Disabled.	
Command Modes	EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

enable secret

Change the password for the `enable` command.

Syntax	`enable secret [level` *`level`*`] [`*`encryption-type`*`] `*`password`* To delete a password, use the `no enable secret [`*`encryption-type`*`] `*`password`*` [level `*`level`*`]` command.	
Parameters	**level** *level*	(OPTIONAL) Enter the keyword `level` then a number as the level of access. The range is from 1 to 15.
	encryption-type	(OPTIONAL) Enter the number `5` or `0` as the encryption type. Enter a `5` then a text string as the hidden password. The text string must be a password that was already encrypted by a Dell Networking router. Use this parameter only with a password that you copied from the `show running-config` file of another Dell Networking router.

	password	Enter a text string, up to 32 characters long, as the clear text password.

Defaults No password is configured. *level* = **15**.

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Usage Information To control access to command modes, use this command to define a password for a level and use the `privilege level (CONFIGURATION mode)` command.

Passwords must meet the following criteria:

- Start with a letter, not a number.
- Passwords can have a regular expression as the password. To create a password with a regular expression in it, use CNTL + v prior to entering regular expression. For example, to create the password `abcd]e`, you type "`abcd CNTL v ]e`". When the password is created, you do not use the CNTL + v key combination and enter "`abcd]e`".

NOTE: The question mark (?) and the tilde (~) are not supported characters.

Related Commands show running-config — views the current configuration.

privilege level (CONFIGURATION mode) — controls access to the command modes within the switch.

exec-banner

Enable the display of a text string when the user enters EXEC mode.

Syntax `exec-banner`

To disable the banner on terminal lines, use the `no exec-banner` command.

Defaults **Enabled on all lines** (if configured, the banner appears).

Command Modes LINE

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Usage Information Optionally, use the `banner exec` command to create a text string that is displayed when you access EXEC mode. This command toggles that display.

Related Commands banner exec — configures a banner to display when entering EXEC mode.

line — enables and configures console and virtual terminal lines to the system.

ip radius source-interface

Specify an interface's IP address as the source IP address for RADIUS connections.

Syntax

`ip radius source-interface` *interface*

To delete a source interface, use the `no ip radius source-interface` command.

Parameters

interface — Enter the following keywords and slot/port or number information:

- For a 100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For Loopback interfaces, enter the keyword `loopback` then a number from zero (0) to 16838.
- For the Null interface, enter the keywords `null 0`.
- For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128.
- For a ten-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For VLAN interface, enter the keyword `vlan` then a number from 1 to 4094.

Defaults Not configured.

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

ip tacacs source-interface

Specify an interface's IP address as the source IP address for TACACS+ connections.

Syntax

`ip tacacs source-interface` *interface*

To delete a source interface, use the `no ip tacacs source-interface` command.

Parameters

interface — Enter the following keywords and slot/port or number information:

- For a 100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.

- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For Loopback interfaces, enter the keyword `loopback` then a number from zero (0) to 16838.
- For the Null interface, enter the keywords `null 0`.
- For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128.
- For a ten-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For VLAN interface, enter the keyword `vlan` then a number from 1 to 4094.

Defaults	Not configured.	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

login authentication

To designate the terminal lines, apply an authentication method list.

Syntax `login authentication {method-list-name | default}`

To use the local user/password database for login authentication, use the `no login authentication` command.

Parameters

method-list-name	Enter the keywords `method-list-name` to specify that method list, created in the `aaa authentication login` command, to be applied to the designated terminal line.
default	Enter the keyword `default` to specify that the default method list, created in the `aaa authentication login` command, is applied to the terminal line.

Defaults	No authentication is performed on the console lines. Local authentication is performed on the virtual terminal and auxiliary lines.
Command Modes	LINE
Command History	**Version 9.3.0.0** Introduced on the M I/O Aggregator.
Usage Information	If you configure the `aaa authentication login default` command, the `login authentication default` command automatically is applied to all terminal lines.

Related Commands	aaa authentication login — selects the login authentication methods.

motd-banner

Enable a message of the day (MOTD) banner to appear when you log in to the system.

Syntax	`motd-banner` To disable the MOTD banner, use the `no motd-banner` command.
Defaults	Enabled on all lines.
Command Modes	LINE

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

password-attributes

Configure the password attributes (strong password).

Syntax	`password-attributes [min-length` *`number`*`] [max-retry` *`number`*`] [character-restriction [upper` *`number`*`] [lower` *`number`*`] [numeric` *`number`*`] [special-char` *`number`*`]]` To return to the default, use the `no password-attributes [min-length` *`number`*`] [max-retry` *`number`*`] [character-restriction [upper` *`number`*`] [lower` *`number`*`] [numeric` *`number`*`] [special-char` *`number`*`]]` command.

Parameters

min-length *number*	(OPTIONAL) Enter the keywords `min-length` then the number of characters. The range is from 0 to 32 characters.
max-retry *number*	(OPTIONAL) Enter the keywords `max-retry` then the number of maximum password retries. The range is from 0 to 16.
character-restriction	(OPTIONAL) Enter the keywords `character-restriction` to indicate a character restriction for the password.
upper *number*	(OPTIONAL) Enter the keyword `upper` then the upper number. The range is from 0 to 31.
lower *number*	(OPTIONAL) Enter the keyword `lower` then the lower number. The range is from 0 to 31.
numeric *number*	(OPTIONAL) Enter the keyword `numeric` then the numeric number. The range is from 0 to 31.
special-char *number*	(OPTIONAL) Enter the keywords `special-char` then the number of special characters permitted. The range is from 0 to 31.

Defaults	none	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.
Related Commands	password — specifies a password for users on terminal lines.	

privilege level (CONFIGURATION mode)

Change the access or privilege level of one or more commands.

Syntax `privilege` *mode* `{level` *level command* | `reset` *command*`}`

To delete access to a level and command, use the `no privilege mode level level command` command.

Parameters

mode	Enter one of the following keywords as the mode for which you are controlling access: • `configure` for CONFIGURATION mode • `exec` for EXEC mode • `interface` for INTERFACE modes • `line` for LINE mode • `route-map` for ROUTE-MAP mode • `router` for ROUTER OSPF, ROUTER RIP, ROUTER ISIS and ROUTER BGP modes
level *level*	Enter the keyword `level` then a number for the access level. The range is from 0 to 15. Level 1 is EXEC mode and Level 15 allows access to all CLI modes and commands.
reset	Enter the keyword `reset` to return the security level to the default setting.
command	Enter the command's keywords to assign the command to a certain access level. You can enter one or all of the keywords.

Defaults	Not configured.	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator

Usage Information	To define a password for the level to which you are assigning privilege or access, use the `enable password` command.

privilege level (LINE mode)

Change the access level for users on the terminal lines.

Syntax	`privilege level level` To delete access to a terminal line, use the `no privilege level level` command.	
Parameters	**level *level***	Enter the keyword `level` then a number for the access level. The range is from 0 to 15. Level 1 is EXEC mode and Level 15 allows access to all CLI modes.
Defaults	*level* = **15**	
Command Modes	LINE	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator

RADIUS Commands

The following RADIUS commands are supported by FTOS.

radius-server deadtime

Configure a time interval during which non-responsive RADIUS servers to authentication requests are skipped.

Syntax	`radius-server deadtime seconds` To disable this function or return to the default value, use the `no radius-server deadtime` command.	
Parameters	***seconds***	Enter a number of seconds during which non-responsive RADIUS servers are skipped. The range is from 0 to 2147483647 seconds. The default is **0 seconds**.
Defaults	**0 seconds**	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

radius-server host

Configure a RADIUS server host.

Syntax

```
radius-server host {hostname | ipv4-address | ipv6-address}
[auth-port port-number] [retransmit retries] [timeout seconds]
[key [encryption-type] key]
```

Parameters

hostname	Enter the name of the RADIUS server host.
ipv4-address | *ipv6-address*	Enter the IPv4 address (A.B.C.D) or IPv6 address (X:X:X:X::X) of the RADIUS server host.
auth-port *port-number*	(OPTIONAL) Enter the keywords `auth-port` then a number as the port number. The range is from zero (0) to 65535. The default port-number is **1812**.
retransmit *retries*	(OPTIONAL) Enter the keyword `retransmit` then a number as the number of attempts. This parameter overwrites the `radius-server retransmit` command. The range is from zero (0) to 100. The default is **3 attempts**.
timeout *seconds*	(OPTIONAL) Enter the keyword `timeout` then the seconds the time interval the switch waits for a reply from the RADIUS server. This parameter overwrites the `radius-server timeout` command. The range is from 0 to 1000. The default is **5 seconds**.
key [*encryption-type*] *key*	(OPTIONAL) Enter the keyword `key` then an optional encryption-type and a string up to 42 characters long as the authentication key. The RADIUS host server uses this authentication key and the RADIUS daemon operating on this switch. For the encryption-type, enter either zero (`0`) or `7` as the encryption type for the key entered. The options are: • `0` is the default and means the password is not encrypted and stored as clear text. • `7` means that the password is encrypted and hidden. Configure this parameter last because leading spaces are ignored.

Defaults Not configured.

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Usage Information	To configure any number of RADIUS server hosts for each server host that is configured, use this command. FTOS searches for the RADIUS hosts in the order they are configured in the software. The global default values for the `timeout`, `retransmit`, and `key` optional parameters are applied, unless those values are specified in the `radius-server host` or other commands. To return to the global default values, if you configure the `timeout`, `retransmit`, or `key` values, include those keywords when using the `no radius-server host` command syntax.
Related Commands	login authentication — sets the database to be checked when a user logs in. radius-server key — sets an authentication key for RADIUS communications. radius-server retransmit — sets the number of times the RADIUS server attempts to send information. radius-server timeout — sets the time interval before the RADIUS server times out.

radius-server retransmit

Configure the number of times the switch attempts to connect with the configured RADIUS host server before declaring the RADIUS host server unreachable.

Syntax	`radius-server retransmit` *`retries`* To configure zero retransmit attempts, use the `no radius-server retransmit` command. To return to the default setting, use the `radius-server retransmit 3` command.	
Parameters	***retries***	Enter a number of attempts that FTOS tries to locate a RADIUS server. The range is from zero (0) to 100. The default is **3 retries**.
Defaults	**3 retries**	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.
Related Commands	radius-server host — configures a RADIUS host.	

radius-server timeout

To reply to a request, configure the amount of time the RADIUS client (the switch) waits for a RADIUS host server .

Syntax	`radius-server timeout` *`seconds`* To return to the default value, use the `no radius-server timeout` command.

Parameters	*seconds*	Enter the number of seconds between an unsuccessful attempt and the FTOS times out. The range is from zero (0) to 1000 seconds. The default is **5 seconds**.
Defaults	**5 seconds**	
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.
Related Commands	radius-server host — configures a RADIUS host.	

radius-server key

Configure a key for all RADIUS communications between the switch and the RADIUS host server.

Syntax

`radius-server key [`*`encryption-type`*`] `*`key`*

To delete a password, use the `no radius-server key` command.

Parameters

encryption-type	(OPTIONAL) Enter either zero (`0`) or `7` as the encryption type for the key entered. The options are: • `0` is the default and means the key is not encrypted and stored as clear text. • `7` means that the key is encrypted and hidden.
key	Enter a string that is the key to be exchanged between the switch and RADIUS servers. It can be up to 42 characters long.

Defaults

Not configured.

Command Modes

CONFIGURATION

Command History	**Version 9.3.0.0** Introduced on the M I/O Aggregator.

Usage Information

The key configured on the switch must match the key configured on the RADIUS server daemon.

If you configure the `key` parameter in the `radius-server host` command, the key configured with the `radius-server key` command is the default key for all RADIUS communications.

Related Commands

radius-server host — configures a RADIUS host.

show privilege

View your access level.

Syntax	`show privilege`	
Command Modes	• EXEC • EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.

Example	`FTOS#show privilege`
	`Current privilege level is 15`
	`FTOS#`	

Related Commands	privilege level (CONFIGURATION mode) — assigns access control to different command modes.	

Suppressing AAA Accounting for Null Username Sessions

When you activate AAA accounting, the FTOS software issues accounting records for all users on the system, including users whose username string, because of protocol translation, is NULL.

An example of this is a user who comes in on a line where the AAA authentication `login method-list none` command is applied. To prevent accounting records from being generated for sessions that do not have usernames associated with them, use the following command.

- Prevent accounting records from being generated for users whose username string is NULL.
  CONFIGURATION mode
  `aaa accounting suppress null-username`

TACACS+ Commands

FTOS supports TACACS+ as an alternate method for login authentication.

tacacs-server host

Specify a TACACS+ host.

Syntax	`tacacs-server host {`*`hostname`* `|` *`ipv4-address`* `|` *`ipv6-address`*`}` `[port` *`number`*`] [timeout` *`seconds`*`] [key` *`key`*`]`	
Parameters	***hostname***	Enter the name of the TACACS+ server host.
	ipv4-address **|** ***ipv6-address***	Enter the IPv4 address (A.B.C.D) or IPv6 address (X:X:X:X::X) of the TACACS+ server host.
	port ***number***	(OPTIONAL) Enter the keyword `port` then a number as the port to be used by the TACACS+ server. The range is from zero (0) to 65535. The default is **49**.

	timeout *seconds*	(OPTIONAL) Enter the keyword `timeout` then the number of seconds the switch waits for a reply from the TACACS+ server. The range is from 0 to 1000. The default is **10 seconds**.
	key *key*	(OPTIONAL) Enter the keyword `key` then a string up to 42 characters long as the authentication key. This authentication key must match the key specified in the `tacacs-server key` for the TACACS+ daemon.

Defaults Not configured.

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Usage Information

To list multiple TACACS+ servers to be used by the `aaa authentication login` command, configure this command multiple times.

If you are not configuring the switch as a TACACS+ server, you do not need to configure the `port`, `timeout` and `key` optional parameters. If you do not configure a key, the key assigned in the `tacacs-server key` command is used.

Related Commands

aaa authentication login — specifies the login authentication method.

tacacs-server key — configures a TACACS+ key for the TACACS server.

tacacs-server key

Configure a key for communication between a TACACS+ server and a client.

Syntax

```
tacacs-server key [encryption-type] key
```

To delete a key, use the `no tacacs-server key key` command.

Parameters

encryption-type	(OPTIONAL) Enter either zero (`0`) or `7` as the encryption type for the key entered. The options are: • `0` is the default and means the key is not encrypted and stored as clear text. • `7` means that the key is encrypted and hidden.
key	Enter a text string, up to 42 characters long, as the clear text password. Leading spaces are ignored.

Defaults Not configured.

Command Modes CONFIGURATION

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.
Usage Information	The key configured with this command must match the key configured on the TACACS+ daemon.	

timeout login response

Specify how long the software waits for the login input (for example, the user name and password) before timing out.

Syntax

`timeout login response` *seconds*

To return to the default values, use the `no timeout login response` command.

Parameters

seconds — Enter a number of seconds the software waits before logging you out. The range is:

- VTY: the range is from 1 to 30 seconds, the default is **30 seconds**.
- Console: the range is from 1 to 300 seconds, the default is **0 seconds** (no timeout).
- AUX: the range is from 1 to 300 seconds, the default is **0 seconds** (no timeout).

Defaults	See the defaults settings shown in *Parameters*.	
Command Modes	LINE	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator.
Usage Information	The software measures the period of inactivity defined in this command as the period between consecutive keystrokes. For example, if your password is "password" you can enter "p" and wait 29 seconds to enter the next letter.	

Understanding Banner Settings

This functionality is supported on the M I/O Aggregator.

A banner is a descriptive, meaningful quote or an informational-note that is displayed when you log in to the system, depending on the privilege level and the command mode into which the user logs in. You can specify different banners to be displayed as the message-of-the-day (MOTD), as the opening quote in EXEC mode, or as the beginning message in EXEC Privilege mode. Setting up a banner to be displayed enables you to display in an easily-noticeable, prominent manner of any important information or welcome, group-level notification that needs to be communicated to all of the users of the system.

A login banner message is displayed only in EXEC Privilege mode after entering the enable command followed by the password. These banners are not displayed to users in EXEC mode. When you connect to a system, the message-of-the-day (MOTD) banner is displayed first, followed by the login banner and prompts. After you log in to the system with valid authentication credentials, the EXEC banner is shown.

You can use the MOTD banner to indicate to users of critical upcoming events such as a lab shutdown of devices or of any circuit-level maintenance or downtime that is upcoming so that they can plan and schedule their accessibility to the device , based on the network outages and system reboots. You can modify the banner messages depending on the requirements or conditions.

AAA Authentication

FTOS supports a distributed client/server system implemented through authentication, authorization, and accounting (AAA) to help secure networks against unauthorized access.

In the Dell Networking implementation, the Dell Networking system acts as a RADIUS or TACACS+ client and sends authentication requests to a central remote authentication dial-in service (RADIUS) or Terminal access controller access control system plus (TACACS+) server that contains all user authentication and network service access information.

Dell Networking uses local usernames/passwords (stored on the Dell Networking system) or AAA for login authentication. With AAA, you can specify the security protocol or mechanism for different login methods and different users. In FTOS, AAA uses a list of authentication methods, called method lists, to define the types of authentication and the sequence in which they are applied. You can define a method list or use the default method list. User-defined method lists take precedence over the default method list.

NOTE: If a console user logs in with RADIUS authentication, the privilege level is applied from the RADIUS server if the privilege level is configured for that user in RADIUS, whether you configure RADIUS authorization.

Configuration Task List for AAA Authentication

The following sections provide the configuration tasks.

- Configure Login Authentication for Terminal Lines
- Configuring AAA Authentication Login Methods
- Enabling AAA Authentication
- Enabling AAA Authentication—RADIUS

For a complete list of all commands related to login authentication, refer to the *Security* chapter in the *FTOS Command Reference Guide*.

Configure Login Authentication for Terminal Lines

You can assign up to five authentication methods to a method list. FTOS evaluates the methods in the order in which you enter them in each list.

If the first method list does not respond or returns an error, FTOS applies the next method list until the user either passes or fails the authentication. If the user fails a method list, FTOS does not apply the next method list.

Configuring AAA Authentication Login Methods

To configure an authentication method and method list, use the following commands.

FTOS Behavior: If you use a method list on the console port in which RADIUS or TACACS is the last authentication method, and the server is not reachable, FTOS allows access even though the username

and password credentials cannot be verified. Only the console port behaves this way, and does so to ensure that users are not locked out of the system if network-wide issue prevents access to these servers.

1. Define an authentication method-list (*method-list-name*) or specify the default.
   CONFIGURATION mode
   ```
   aaa authentication login {method-list-name | default} method1 [... method4]
   ```
   The default method-list is applied to all terminal lines.
   Possible methods are:
   - `enable`: use the password you defined using the `enable secret` or `enable password` command in CONFIGURATION mode.
   - `line`: use the password you defined using the `password` command in LINE mode.
   - `local`: use the username/password database defined in the local configuration.
   - `none`: no authentication.
   - `radius`: use the RADIUS servers configured with the radius-server host command.
   - `tacacs+`: use the TACACS+ servers configured with the `tacacs-server host` command.
2. Enter LINE mode.
   CONFIGURATION mode
   ```
   line {aux 0 | console 0 | vty number [... end-number]}
   ```
3. Assign a *method-list-name* or the default list to the terminal line.
   LINE mode
   ```
   login authentication {method-list-name | default}
   ```

To view the configuration, use the `show config` command in LINE mode or the `show running-config` in EXEC Privilege mode.

NOTE: Dell Networking recommends using the `none` method only as a backup. This method does not authenticate users. The `none` and `enable` methods do not work with secure shell (SSH).

You can create multiple method lists and assign them to different terminal lines.

Enabling AAA Authentication

To enable AAA authentication, use the following command.

- Enable AAA authentication.
  CONFIGURATION mode
  ```
  aaa authentication enable {method-list-name | default} method1 [... method4]
  ```
  - `default`: uses the listed authentication methods that follow this argument as the default list of methods when a user logs in.
  - *method-list-name*: character string used to name the list of enable authentication methods activated when a user logs in.
  - *method1 [... method4]*: any of the following: RADIUS, TACACS, enable, line, none.

If you do not set the default list, only the local enable is checked. This setting has the same effect as issuing an `aaa authentication enable default enable` command.

Enabling AAA Authentication — RADIUS

To enable authentication from the RADIUS server, and use TACACS as a backup, use the following commands.

1. Enable RADIUS and set up TACACS as backup.
   CONFIGURATION mode
   ```
   aaa authentication enable default radius tacacs
   ```
2. Establish a host address and password.
   CONFIGURATION mode
   ```
   radius-server host x.x.x.x key some-password
   ```
3. Establish a host address and password.
   CONFIGURATION mode
   ```
   tacacs-server host x.x.x.x key some-password
   ```

To get `enable authentication` from the RADIUS server and use TACACS as a backup, issue the following commands.

Example of Enabling Authentication from the RADIUS Server

```
FTOS(config)# aaa authentication enable default radius tacacs
Radius and TACACS server has to be properly setup for this.
FTOS(config)# radius-server host x.x.x.x key <some-password>
FTOS(config)# tacacs-server host x.x.x.x key <some-password>
```

To use local authentication for `enable secret` on the console, while using remote authentication on VTY lines, issue the following commands.

Example of Enabling Local Authentication for the Console and Remote Authentication for VTY Lines

```
FTOS(config)# aaa authentication enable mymethodlist radius tacacs
FTOS(config)# line vty 0 9
FTOS(config-line-vty)# enable authentication mymethodlist
```

Server-Side Configuration

- **TACACS+** — When using TACACS+, Dell Networking sends an initial packet with service type SVC_ENABLE, and then sends a second packet with just the password. The TACACS server must have an entry for username $enable$.
- **RADIUS** — When using RADIUS authentication, FTOS sends an authentication packet with the following:
  ```
  Username: $enab15$
  Password: <password-entered-by-user>
  ```

Therefore, the RADIUS server must have an entry for this username.

RADIUS

Remote authentication dial-in user service (RADIUS) is a distributed client/server protocol.

This protocol transmits authentication, authorization, and configuration information between a central RADIUS server and a RADIUS client (the Dell Networking system). The system sends user information to the RADIUS server and requests authentication of the user and password. The RADIUS server returns one of the following responses:

- **Access-Accept** — the RADIUS server authenticates the user.

- **Access-Reject** — the RADIUS server does not authenticate the user.

If an error occurs in the transmission or reception of RADIUS packets, you can view the error by enabling the `debug radius` command.

Transactions between the RADIUS server and the client are encrypted (the users' passwords are not sent in plain text). RADIUS uses UDP as the transport protocol between the RADIUS server host and the client.

For more information about RADIUS, refer to RFC 2865, *Remote Authentication Dial-in User Service*.

RADIUS Authentication and Authorization

FTOS supports RADIUS for user authentication (text password) at login and can be specified as one of the login authentication methods in the `aaa authentication login` command.

When configuring AAA authorization, you can configure to limit the attributes of services available to a user. When you enable authorization, the network access server uses configuration information from the user profile to issue the user's session. The user's access is limited based on the configuration attributes.

RADIUS exec-authorization stores a user-shell profile and that is applied during user login. You may name the relevant named-lists with either a unique name or the default name. When you enable authorization by the RADIUS server, the server returns the following information to the client:

- Idle Time
- ACL Configuration Information
- Auto-Command
- Privilege Levels

After gaining authorization for the first time, you may configure these attributes.

NOTE: RADIUS authentication/authorization is done for every login. There is no difference between first-time login and subsequent logins.

Idle Time

Every session line has its own idle-time. If the idle-time value is not changed, the default value of **30 minutes** is used.

RADIUS specifies idle-time allow for a user during a session before timeout. When a user logs in, the lower of the two idle-time values (configured or default) is used. The idle-time value is updated if both of the following happens:

- The administrator changes the idle-time of the line on which the user has logged in.
- The idle-time is lower than the RADIUS-returned idle-time.

ACL Configuration Information

The RADIUS server can specify an ACL. If an ACL is configured on the RADIUS server, and if that ACL is present, the user may be allowed access based on that ACL.

If the ACL is absent, authorization fails, and a message is logged indicating this.

RADIUS can specify an ACL for the user if both of the following are true:

- If an ACL is absent.
- If there is a very long delay for an entry, or a denied entry because of an ACL, and a message is logged.

NOTE: The ACL name must be a string. Only standard ACLs in authorization (both RADIUS and TACACS) are supported. Authorization is denied in cases using Extended ACLs.

Auto-Command

You can configure the system through the RADIUS server to automatically execute a command when you connect to a specific line.

The `auto-command` command is executed when the user is authenticated and before the prompt appears to the user.

- Automatically execute a command.

  `auto-command`

Privilege Levels

Through the RADIUS server, you can configure a privilege level for the user to enter into when they connect to a session.

This value is configured on the client system.

- Set a privilege level.

  `privilege level`

Configuration Task List for RADIUS

To authenticate users using RADIUS, you must specify at least one RADIUS server so that the system can communicate with and configure RADIUS as one of your authentication methods.

The following list includes the configuration tasks for RADIUS.

- Defining a AAA Method List to be Used for RADIUS (mandatory)
- Applying the Method List to Terminal Lines (mandatory except when using default lists)
- Specifying a RADIUS Server Host (mandatory)
- Setting Global Communication Parameters for all RADIUS Server Hosts (optional)
- Monitoring RADIUS (optional)

For a complete listing of all FTOS commands related to RADIUS, refer to the *Security* chapter in the *FTOS Command Reference Guide*.

NOTE: RADIUS authentication and authorization are done in a single step. Hence, authorization cannot be used independent of authentication. However, if you have configured RADIUS authorization and have not configured authentication, a message is logged stating this. During authorization, the next method in the list (if present) is used, or if another method is not present, an error is reported.

To view the configuration, use the `show config` in LINE mode or the `show running-config` command in EXEC Privilege mode.

Defining a AAA Method List to be Used for RADIUS

To configure RADIUS to authenticate or authorize users on the system, create a AAA method list.

Default method lists do not need to be explicitly applied to the line, so they are not mandatory.

To create a method list, use the following commands.

- Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the RADIUS authentication method.

  CONFIGURATION mode

  `aaa authentication login` *`method-list-name`* `radius`

- Create a method list with RADIUS and TACACS+ as authorization methods.

  CONFIGURATION mode

  ```
  aaa authorization exec {method-list-name | default} radius tacacs+
  ```

  Typical order of methods: RADIUS, TACACS+, Local, None.

  If RADIUS denies authorization, the session ends (RADIUS must not be the last method specified).

Applying the Method List to Terminal Lines

To enable RADIUS AAA login authentication for a method list, apply it to a terminal line.

To configure a terminal line for RADIUS authentication and authorization, use the following commands.

- Enter LINE mode.

  CONFIGURATION mode

  ```
  line {aux 0 | console 0 | vty number [end-number]}
  ```

- Enable AAA login authentication for the specified RADIUS method list.

  LINE mode

  ```
  login authentication {method-list-name | default}
  ```

  This procedure is mandatory if you are not using default lists.

- To use the method list.

  CONFIGURATION mode

  ```
  authorization exec methodlist
  ```

Specifying a RADIUS Server Host

When configuring a RADIUS server host, you can set different communication parameters, such as the UDP port, the key password, the number of retries, and the timeout.

To specify a RADIUS server host and configure its communication parameters, use the following command.

- Enter the host name or IP address of the RADIUS server host.

  CONFIGURATION mode

  ```
  radius-server host {hostname | ip-address} [auth-port port-number]
  [retransmit retries] [timeout seconds] [key [encryption-type] key]
  ```

  Configure the optional communication parameters for the specific host:

  - `auth-port port-number`: the range is from 0 to 65335. Enter a UDP port number. The default is **1812**.
  - `retransmit retries`: the range is from 0 to 100. Default is **3**.
  - `timeout seconds`: the range is from 0 to 1000. Default is **5 seconds**.
  - `key [encryption-type] key`: enter 0 for plain text or 7 for encrypted text, and a string for the key. The key can be up to 42 characters long. This key must match the key configured on the RADIUS server host.

  If you do not configure these optional parameters, the global default values for all RADIUS host are applied.

To specify multiple RADIUS server hosts, configure the `radius-server host` command multiple times. If you configure multiple RADIUS server hosts, FTOS attempts to connect with them in the order in which they were configured. When FTOS attempts to authenticate a user, the software connects with the RADIUS server hosts one at a time, until a RADIUS server host responds with an accept or reject response.

If you want to change an optional parameter setting for a specific host, use the `radius-server host` command. To change the global communication settings to all RADIUS server hosts, refer to Setting Global Communication Parameters for all RADIUS Server Hosts.

To view the RADIUS configuration, use the `show running-config radius` command in EXEC Privilege mode.

To delete a RADIUS server host, use the `no radius-server host {hostname | ip-address}` command.

Setting Global Communication Parameters for all RADIUS Server Hosts

You can configure global communication parameters (`auth-port`, `key`, `retransmit`, and `timeout` parameters) and specific host communication parameters on the same system.

However, if you configure both global and specific host parameters, the specific host parameters override the global parameters for that RADIUS server host.

To set global communication parameters for all RADIUS server hosts, use the following commands.

- Set a time interval after which a RADIUS host server is declared dead.

  CONFIGURATION mode

  `radius-server deadtime seconds`

  - *seconds*: the range is from 0 to 2147483647. The default is **0 seconds**.
- Configure a key for all RADIUS communications between the system and RADIUS server hosts.

  CONFIGURATION mode

  `radius-server key [encryption-type] key`

  - *encryption-type*: enter 7 to encrypt the password. Enter 0 to keep the password as plain text.
  - *key*: enter a string. The key can be up to 42 characters long. You cannot use spaces in the key.
- Configure the number of times FTOS retransmits RADIUS requests.

  CONFIGURATION mode

  `radius-server retransmit retries`

  - *retries*: the range is from 0 to 100. Default is **3 retries**.
- Configure the time interval the system waits for a RADIUS server host response.

  CONFIGURATION mode

  `radius-server timeout seconds`

  - *seconds*: the range is from 0 to 1000. Default is **5 seconds**.

To view the configuration of RADIUS communication parameters, use the `show running-config` command in EXEC Privilege mode.

Monitoring RADIUS

To view information on RADIUS transactions, use the following command.

- View RADIUS transactions to troubleshoot problems.

  EXEC Privilege mode

  `debug radius`

TACACS+

FTOS supports terminal access controller access control system (TACACS+ client, including support for login authentication.

Configuration Task List for TACACS+

The following list includes the configuration task for TACACS+ functions.

- Choosing TACACS+ as the Authentication Method
- Monitoring TACACS+
- TACACS+ Remote Authentication and Authorization
- Specifying a TACACS+ Server Host

For a complete listing of all commands related to TACACS+, refer to the *Security* chapter in the *FTOS Command Reference Guide*.

Choosing TACACS+ as the Authentication Method

One of the login authentication methods available is TACACS+ and the user's name and password are sent for authentication to the TACACS hosts specified.

To use TACACS+ to authenticate users, specify at least one TACACS+ server for the system to communicate with and configure TACACS+ as one of your authentication methods.

To select TACACS+ as the login authentication method, use the following commands.

1. Configure a TACACS+ server host.

   CONFIGURATION mode

   ```
   tacacs-server host {ip-address | host}
   ```

   Enter the IP address or host name of the TACACS+ server.

   Use this command multiple times to configure multiple TACACS+ server hosts.
2. Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the TACAS+ authentication method.

   CONFIGURATION mode

   ```
   aaa authentication login {method-list-name | default} tacacs+ [...method3]
   ```

   The TACACS+ method must not be the last method specified.
3. Enter LINE mode.

   CONFIGURATION mode

   ```
   line {aux 0 | console 0 | vty number [end-number]}
   ```
4. Assign the *method-list* to the terminal line.

   LINE mode

   ```
   login authentication {method-list-name | default}
   ```

To view the configuration, use the `show config` in LINE mode or the `show running-config tacacs+` command in EXEC Privilege mode.

If authentication fails using the primary method, FTOS employs the second method (or third method, if necessary) automatically. For example, if the TACACS+ server is reachable, but the server key is invalid, FTOS proceeds to the next authentication method. In the following example, the TACACS+ is incorrect, but the user is still authenticated by the secondary method.

First bold line: Server key purposely changed to incorrect value.

Second bold line: User authenticated using the secondary method.

Example of a Failed Authentication

```
FTOS(conf)#
FTOS(conf)#do show run aaa
!
aaa authentication enable default tacacs+ enable
aaa authentication enable LOCAL enable tacacs+
aaa authentication login default tacacs+ local
aaa authentication login LOCAL local tacacs+
aaa authorization exec default tacacs+ none
aaa authorization commands 1 default tacacs+ none
aaa authorization commands 15 default tacacs+ none
aaa accounting exec default start-stop tacacs+
aaa accounting commands 1 default start-stop tacacs+
aaa accounting commands 15 default start-stop tacacs+
FTOS(conf)#
FTOS(conf)#do show run tacacs+
!
tacacs-server key 7 d05206c308f4d35b
tacacs-server host 10.10.10.10 timeout 1
FTOS(conf)#tacacs-server key angeline
FTOS(conf)#%RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user admin on
vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password
authentication success on vty0 ( 10.11.9.209 )
%RPM0-P:CP %SEC-5-LOGOUT: Exec session is terminated for user admin on line
vty0 (10.11.9.209)
FTOS(conf)#username angeline password angeline
FTOS(conf)#%RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user angeline
on vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password
authentication success on vty0 ( 10.11.9.209 )
```

Monitoring TACACS+

To view information on TACACS+ transactions, use the following command.

- View TACACS+ transactions to troubleshoot problems.

  EXEC Privilege mode

  ```
  debug tacacs+
  ```

TACACS+ Remote Authentication and Authorization

FTOS takes the access class from the TACACS+ server. Access class is the class of service that restricts Telnet access and packet sizes.

If you have configured remote authorization, FTOS ignores the access class you have configured for the VTY line. FTOS instead gets this access class information from the TACACS+ server. FTOS must know the username and password of the incoming user before it can fetch the access class from the server. A user, therefore, at least sees the login prompt. If the access class denies the connection, FTOS closes the Telnet session immediately.

The following example demonstrates how to configure the access-class from a TACACS+ server. This configuration ignores the configured access-class on the VTY line. If you have configured a deny10 ACL on the TACACS+ server, FTOS downloads it and applies it. If the user is found to be coming from the 10.0.0.0 subnet, FTOS also immediately closes the Telnet connection. Note, that no matter where the user is coming from, they see the login prompt.

When configuring a TACACS+ server host, you can set different communication parameters, such as the key password.

Example of Specifying a TACACS+ Server Host

```
FTOS#
FTOS(conf)#
FTOS(conf)#ip access-list standard deny10
FTOS(conf-std-nacl)#permit 10.0.0.0/8
FTOS(conf-std-nacl)#deny any
FTOS(conf)#
FTOS(conf)#aaa authentication login tacacsmethod tacacs+
FTOS(conf)#aaa authentication exec tacacsauthorization tacacs+
FTOS(conf)#tacacs-server host 25.1.1.2 key Force10
FTOS(conf)#
FTOS(conf)#line vty 0 9
FTOS(config-line-vty)#login authentication tacacsmethod
FTOS(config-line-vty)#authorization exec tacauthor
FTOS(config-line-vty)#
FTOS(config-line-vty)#access-class deny10
FTOS(config-line-vty)#end
```

Specifying a TACACS+ Server Host

To specify a TACACS+ server host and configure its communication parameters, use the following command.

- Enter the host name or IP address of the TACACS+ server host.

  CONFIGURATION mode

  `tacacs-server host {`*`hostname`* `|` *`ip-address`*`} [port` *`port-number`*`] [timeout` *`seconds`*`] [key` *`key`*`]`

  Configure the optional communication parameters for the specific host:

  - `port` *`port-number`*: the range is from 0 to 65335. Enter a TCP port number. The default is **49**.
  - `timeout` *`seconds`*: the range is from 0 to 1000. Default is **10 seconds**.
  - `key` *`key`*: enter a string for the key. The key can be up to 42 characters long. This key must match a key configured on the TACACS+ server host. This parameter must be the last parameter you configure.

  If you do not configure these optional parameters, the default global values are applied.

To specify multiple TACACS+ server hosts, configure the `tacacs-server host` command multiple times. If you configure multiple TACACS+ server hosts, FTOS attempts to connect with them in the order in which they were configured.

To view the TACACS+ configuration, use the `show running-config tacacs+` command in EXEC Privilege mode.

To delete a TACACS+ server host, use the `no tacacs-server host {`*`hostname`* `|` *`ip-address`*`}` command.

Example of Connecting with a TACACS+ Server Host

```
freebsd2# telnet 2200:2200:2200:2200:2200::2202
Trying 2200:2200:2200:2200:2200::2202...
Connected to 2200:2200:2200:2200:2200::2202.
Escape character is '^]'.
Login: admin
Password:
FTOS#
FTOS#
```

Command Authorization

The AAA command authorization feature configures FTOS to send each configuration command to a TACACS server for authorization before it is added to the running configuration.

By default, the AAA authorization commands configure the system to check both EXEC mode and CONFIGURATION mode commands. Use the `no aaa authorization config-commands` command to enable only EXEC mode command checking.

If rejected by the AAA server, the command is not added to the running config, and a message displays:

```
04:07:48: %RPM0-P:CP %SEC-3-SEC_AUTHORIZATION_FAIL: Authorization failure
Command
authorization failed for user (denyall) on vty0 ( 10.11.9.209 )
```

Protection from TCP Tiny and Overlapping Fragment Attacks

Tiny and overlapping fragment attack is a class of attack where configured ACL entries — denying TCP port-specific traffic — is bypassed and traffic is sent to its destination although denied by the ACL.

RFC 1858 and 3128 proposes a countermeasure to the problem. This countermeasure is configured into the line cards and enabled by default.

Enabling SCP and SSH

Secure shell (SSH) is a protocol for secure remote login and other secure network services over an insecure network. FTOS is compatible with SSH versions 1.5 and 2, both the client and server modes. SSH sessions are encrypted and use authentication.

For details about the command syntax, refer to the *Security* chapter in the *FTOS Command Line Interface Reference Guide*.

SCP is a remote file copy program that works with SSH and FTOS supports.

NOTE: The Windows-based WinSCP client software is not supported for secure copying between a PC and an FTOS-based system. Unix-based SCP client software is supported.

To use the SSH client, use the following command.

- Open an SSH connection and specifying the host name, username, port number, and version of the SSH client.

  EXEC Privilege mode

  `ssh {hostname} [-l username | -p port-number | -v {1 | 2}`

  hostname is the IP address or host name of the remote device. Enter an IPv4 or IPv6 address in dotted decimal format (A.B.C.D).
- Configure the Dell Networking system as an SCP/SSH server.

  CONFIGURATION mode

  `ip ssh server {enable | port port-number}`
- Configure the Dell Networking system as an SSH server that uses only version 1 or 2.

  CONFIGURATION mode

  `ip ssh server version {1|2}`
- Display SSH connection information.

  EXEC Privilege mode

```
show ip ssh
```

The following example shows using the `ip ssh server version 2` command to enable SSH version 2 and the `show ip ssh` command to confirm the setting.

Specifying an SSH Version

```
FTOS(conf)#ip ssh server version 2
FTOS(conf)#do show ip ssh
SSH server               : disabled.
SSH server version       : v2.
Password Authentication  : enabled.
Hostbased Authentication : disabled.
RSA Authentication       : disabled.
```

To disable SSH server functions, use the `no ip ssh server enable` command.

Using SCP with SSH to Copy a Software Image

To use secure copy (SCP) to copy a software image through an SSH connection from one switch to another, use the following commands.

1. On Chassis One, set the SSH port number (**port 22** by default).
   CONFIGURATION mode
   `ip ssh server port` *number*
2. On Chassis One, enable SSH.
   CONFIGURATION mode
   `ip ssh server enable`
3. On Chassis Two, invoke SCP.
   CONFIGURATION mode
   `copy scp: flash:`
4. On Chassis Two, in response to prompts, enter the path to the desired file and enter the port number specified in Step 1.
   EXEC Privilege mode

Other SSH-related commands include:

- `crypto key generate`: generate keys for the SSH server.
- `debug ip ssh`: enables collecting SSH debug information.
- `ip scp topdir`: identify a location for files used in secure copy transfer.
- `ip ssh authentication-retries`: configure the maximum number of attempts that should be used to authenticate a user.
- `ip ssh connection-rate-limit`: configure the maximum number of incoming SSH connections per minute.
- `ip ssh hostbased-authentication enable`: enable host-based authentication for the SSHv2 server.
- `ip ssh key-size`: configure the size of the server-generated RSA SSHv1 key.
- `ip ssh password-authentication enable`: enable password authentication for the SSH server.
- `ip ssh pub-key-file`: specify the file the host-based authentication uses.
- `ip ssh rhostsfile`: specify the rhost file the host-based authorization uses.
- `ip ssh rsa-authentication enable`: enable RSA authentication for the SSHv2 server.

- `ip ssh rsa-authentication`: add keys for the RSA authentication.
- `show crypto`: display the public part of the SSH host-keys.
- `show ip ssh client-pub-keys`: display the client public keys used in host-based authentication.
- `show ip ssh rsa-authentication`: display the authorized-keys for the RSA authentication.
- `ssh-peer-rpm`: open an SSH connection to the peer RPM.

The following example shows the use of SCP and SSH to copy a software image from one switch running SSH server on UDP port 99 to the local switch.

Example of Using SCP to Copy from an SSH Server on Another Switch

```
FTOS#copy scp: flash:
Address or name of remote host []: 10.10.10.1
Port number of the server [22]: 99
Source file name []: test.cfg
User name to login remote host: admin
Password to login remote host:
```

Secure Shell Authentication

Secure Shell (SSH) is disabled by default.

Enable SSH using the `ip ssh server enable` command.

SSH supports three methods of authentication:

- Enabling SSH Authentication by Password
- Using RSA Authentication of SSH
- Configuring Host-Based SSH Authentication

Important Points to Remember

- If you enable more than one method, the order in which the methods are preferred is based on the *ssh_config* file on the Unix machine.
- When you enable all the three authentication methods, password authentication is the backup method when the RSA method fails.
- The files *known_hosts* and *known_hosts2* are generated when a user tries to SSH using version 1 or version 2, respectively.

Enabling SSH Authentication by Password

Authenticate an SSH client by prompting for a password when attempting to connect to the Dell Networking system. This setup is the simplest method of authentication and uses SSH version 1.

To enable SSH password authentication, use the following command.

- Enable SSH password authentication.

  CONFIGURATION mode

  `ip ssh password-authentication enable`

To view your SSH configuration, use the `show ip ssh` command from EXEC Privilege mode.

Example of Enabling SSH Password Authentication

```
FTOS(conf)#ip ssh server enable
% Please wait while SSH Daemon initializes ... done.
FTOS(conf)#ip ssh password-authentication enable
FTOS#sh ip ssh
SSH server               : enabled.
Password Authentication  : enabled.
```

```
Hostbased Authentication : disabled.
RSA Authentication       : disabled.
```

Using RSA Authentication of SSH

The following procedure authenticates an SSH client based on an RSA key using RSA authentication. This method uses SSH version 2.

1. On the SSH client (Unix machine), generate an RSA key, as shown in the following example.
2. Copy the public key *id_rsa.pub* to the Dell Networking system.
3. Disable password authentication if enabled.

   CONFIGURATION mode

   ```
   no ip ssh password-authentication enable
   ```
4. Bind the public keys to RSA authentication.

   EXEC Privilege mode

   ```
   ip ssh rsa-authentication enable
   ```
5. Bind the public keys to RSA authentication.

   EXEC Privilege mode

   ```
   ip ssh rsa-authentication my-authorized-keys flash://public_key
   ```

Example of Generating RSA Keys

```
admin@Unix_client#ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/home/admin/.ssh/id_rsa):
/home/admin/.ssh/id_rsa already exists.
Overwrite (y/n)? y
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/admin/.ssh/id_rsa.
Your public key has been saved in /home/admin/.ssh/id_rsa.pub.
```

Configuring Host-Based SSH Authentication

Authenticate a particular host. This method uses SSH version 2.

To configure host-based authentication, use the following commands.

1. Configure RSA Authentication. Refer to Using RSA Authentication of SSH.
2. Create *shosts* by copying the public RSA key to the file *shosts* in the directory *.ssh*, and write the IP address of the host to the file.

   ```
   cp /etc/ssh/ssh_host_rsa_key.pub /.ssh/shosts
   ```

   Refer to the first example.
3. Create a list of IP addresses and usernames that are permitted to SSH in a file called *rhosts*.

   Refer to the second example.
4. Copy the file *shosts* and *rhosts* to the Dell Networking system.
5. Disable password authentication and RSA authentication, if configured

   CONFIGURATION mode or EXEC Privilege mode

   `no ip ssh password-authentication` or `no ip ssh rsa-authentication`
6. Enable host-based authentication.

   CONFIGURATION mode

   ```
   ip ssh hostbased-authentication enable
   ```

7. Bind *shosts* and *rhosts* to host-based authentication.

   CONFIGURATION mode

   `ip ssh pub-key-file flash://filename` or `ip ssh rhostsfile flash://filename`

Example of Creating *shosts*

```
admin@Unix_client# cd /etc/ssh

admin@Unix_client# ls
moduli     sshd_config      ssh_host_dsa_key.pub      ssh_host_key.pub
ssh_host_rsa_key.pub ssh_config ssh_host_dsa_key ssh_host_key
ssh_host_rsa_key

admin@Unix_client# cat ssh_host_rsa_key.pub

ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA8K7jLZRVfjgHJzUOmXxuIbZx/
AyWhVgJDQh39k8v3e8eQvLnHBIsqIL8jVy1QHhUeb7GaDlJVEDAMz30myqQbJgXBBRTWgBpLWwL/
doyUXFufjiL9YmoVTkbKcFmxJEMkE3JyHanEi7hg34LChjk9hL1by8cYZP2kYS2lnSyQWk=

admin@Unix_client# ls

id_rsa id_rsa.pub shosts

admin@Unix_client# cat shosts

10.16.127.201, ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA8K7jLZRVfjgHJzUOmXxuIbZx/AyW
hVgJDQh39k8v3e8eQvLnHBIsqIL8jVy1QHhUeb7GaDlJVEDAMz30myqQbJgXBBRTWgBpLWwL/
doyUXFufjiL9YmoVTkbKcFmxJEMkE3JyHanEi7hg34LChjk9hL1by8cYZP2kYS2lnSyQWk=
```

Example of Creating *rhosts*

```
admin@Unix_client# ls
id_rsa id_rsa.pub rhosts shosts
admin@Unix_client# cat rhosts
10.16.127.201 admin
```

Using Client-Based SSH Authentication

To SSH from the chassis to the SSH client, use the following command.

This method uses SSH version 1 or version 2. If the SSH port is a non-default value, use the `ip ssh server port number` command to change the default port number. You may only change the port number when SSH is disabled. Then use the `-p` option with the `ssh` command.

- SSH from the chassis to the SSH client.

  `ssh ip_address`

Example of Client-Based SSH Authentication

```
FTOS#ssh 10.16.127.201 ?
-l   User name option
-p   SSH server port option (default 22)
-v   SSH protocol version
```

Troubleshooting SSH

To troubleshoot SSH, use the following information.

You may not bind *id_rsa.pub* to RSA authentication while logged in via the console. In this case, this message displays:`%Error: No username set for this term.`

Enable host-based authentication on the server (Dell Networking system) and the client (Unix machine). The following message appears if you attempt to log in via SSH and host-based is disabled on the client.

In this case, verify that host-based authentication is set to "Yes" in the *file ssh_config* (root permission is required to edit this file): `permission denied (host based).`

If the IP address in the RSA key does not match the IP address from which you attempt to log in, the following message appears. In this case, verify that the name and IP address of the client is contained in the *file /etc/*hosts: `RSA Authentication Error`.

Telnet

To use Telnet with SSH, first enable SSH, as previously described.

By default, the Telnet daemon is enabled. If you want to disable the Telnet daemon, use the following command, or disable Telnet in the startup config. To enable or disable the Telnet daemon, use the `[no] ip telnet server enable` command.

Example of Using Telnet for Remote Login

```
FTOS(conf)#ip telnet server enable
FTOS(conf)#no ip telnet server enable
```

VTY Line and Access-Class Configuration

Various methods are available to restrict VTY access in FTOS. These depend on which authentication scheme you use — line, local, or remote.

Table 8. VTY Access

Authentication Method	VTY access-class support?	Username access-class support?	Remote authorization support?
Line	YES	NO	NO
Local	NO	YES	NO
TACACS+	YES	NO	YES (with FTOS version 5.2.1.0 and later)
RADIUS	YES	NO	YES (with FTOS version 6.1.1.0 and later)

FTOS provides several ways to configure access classes for VTY lines, including:

- VTY Line Local Authentication and Authorization
- VTY Line Remote Authentication and Authorization

VTY Line Local Authentication and Authorization

FTOS retrieves the access class from the local database.

To use this feature:

1. Create a username.
2. Enter a password.
3. Assign an access class.
4. Enter a privilege level.

You can assign line authentication on a per-VTY basis; it is a simple password authentication, using an access-class as authorization.

Configure local authentication globally and configure access classes on a per-user basis.

FTOS can assign different access classes to different users by username. Until users attempt to log in, FTOS does not know if they will be assigned a VTY line. This means that incoming users always see a login prompt even if you have excluded them from the VTY line with a deny-all access class. After users identify themselves, FTOS retrieves the access class from the local database and applies it. (FTOS then can close the connection if a user is denied access.)

NOTE: If a VTY user logs in with RADIUS authentication, the privilege level is applied from the RADIUS server only if you configure RADIUS authentication.

The following example shows how to allow or deny a Telnet connection to a user. Users see a login prompt even if they cannot log in. No access class is configured for the VTY line. It defaults from the local database.

NOTE: For more information, refer to Access Control Lists (ACLs).

Example of Configuring VTY Authorization Based on Access Class Retrieved from a Local Database (Per User)

```
FTOS(conf)#user gooduser password abc privilege 10 access-class permitall
FTOS(conf)#user baduser password abc privilege 10 access-class denyall
FTOS(conf)#
FTOS(conf)#aaa authentication login localmethod local
FTOS(conf)#
FTOS(conf)#line vty 0 9
FTOS(config-line-vty)#login authentication localmethod
FTOS(config-line-vty)#end
```

VTY Line Remote Authentication and Authorization

FTOS retrieves the access class from the VTY line.

The Dell Networking operating system (FTOS) takes the access class from the VTY line and applies it to ALL users. FTOS does not need to know the identity of the incoming user and can immediately apply the access class. If the authentication method is RADIUS, TACACS+, or line, and you have configured an access class for the VTY line, FTOS immediately applies it. If the access-class is set to deny all or deny for the incoming subnet, FTOS closes the connection without displaying the login prompt. The following example shows how to deny incoming connections from subnet 10.0.0.0 without displaying a login prompt. The example uses TACACS+ as the authentication mechanism.

Example of Configuring VTY Authorization Based on Access Class Retrieved from the Line (Per Network Address)

```
FTOS(conf)#ip access-list standard deny10
FTOS(conf-ext-nacl)#permit 10.0.0.0/8
FTOS(conf-ext-nacl)#deny any
FTOS(conf)#
FTOS(conf)#aaa authentication login tacacsmethod tacacs+
FTOS(conf)#tacacs-server host 256.1.1.2 key Force10
FTOS(conf)#
FTOS(conf)#line vty 0 9
FTOS(config-line-vty)#login authentication tacacsmethod
FTOS(config-line-vty)#
FTOS(config-line-vty)#access-class deny10
FTOS(config-line-vty)#end
(same applies for radius and line authentication)
```

VTY MAC-SA Filter Support

FTOS supports MAC access lists which permit or deny users based on their source MAC address.

With this approach, you can implement a security policy based on the source MAC address.

To apply a MAC ACL on a VTY line, use the same `access-class` command as IP ACLs.

The following example shows how to deny incoming connections from subnet 10.0.0.0 without displaying a login prompt.

Example of Configuring VTY Authorization Based on MAC ACL for the Line (Per MAC Address)

```
FTOS(conf)#mac access-list standard sourcemac
FTOS(config-std-mac)#permit 00:00:5e:00:01:01
FTOS(config-std-mac)#deny any
FTOS(conf)#
FTOS(conf)#line vty 0 9
FTOS(config-line-vty)#access-class sourcemac
FTOS(config-line-vty)#end
```

17

Simple Network Management Protocol (SNMP)

This chapter describes the SNMP enhancements and contains the following sections:

- FIPS Compatibility Support for SNMPv3

SNMPv3 Compliance With FIPS

This functionality is supported on the S4810, S4820T, S6000, Z9000, I/O Aggregator, and MXL platforms.

SNMPv3 is compliant with the Federal information processing standard (FIPS) cryptography standard. The Advanced Encryption Standard (AES) Cipher Feedback (CFB) 128-bit encryption algorithm is in compliance with RFC 3826. SNMPv3 provides multiple authentication and privacy options for user configuration. A subset of these are FIPS-approved algorithms: HMAC-SHA1-96 for authentication and AES128-CFB for privacy. The other options are not FIPS-approved algorithms because of known security weaknesses. Starting with Dell Networking OS Release 9.3(0.0), the AES128-CFB privacy option is supported and it is compliant with RFC 3826.

Starting with the Dell Networking OS Release 9.3.0.0, the SNMPv3 feature also uses a FIPS-validated cryptographic module for all of its cryptographic operations when the system is configured with the `fips mode enable` in Global Configuration mode. When FIPS mode is enabled on the system, SNMPv3 operates in a FIPS-compliant manner, and only the FIPS-approved algorithm options are available for SNMPv3 user configuration. When FIPS mode is disabled on the system, all options are available for SNMPv3 user configuration.

The following table describes the authentication and privacy options that can be configured when FIPS mode is enabled or disabled:

FIPS Mode	Privacy Options	Authentication Options
Disabled	des56 (DES56-CBC) aes128 (AES128-CFB)	md5 (HMAC-MD5-96) sha (HMAC-SHA1-96)
Enabled	aes128 (AES128-CFB)	sha (HMAC-SHA1-96)

To enable robust, effective protection and security for SNMP packets transferred between the server and the client, you can use the `snmp-server user` *`username`* `group` *`groupname`* `3 auth` *`authentication-type auth-password`* `priv aes128` *`priv-password`* command to specify that AES-CFB 128 encryption algorithm needs to be used.

```
Dell(conf)#snmp-server user snmpguy snmpmon 3 auth sha AArt61wq priv aes128
jntRR59a
```

In this example, for a specified user and a group, the AES128-CFB algorithm is configured, and the authentication password to enable the server receive packets from the host and the privacy password to encode the message contents are configured.

SHA authentication needs to be used with the AES-CFB128 privacy algorithm only when FIPS is enabled because SHA is then the only available authentication level. If FIPS is disabled, you can use MD5 authentication in addition to SHA authentication with the AES-CFB128 privacy algorithm

You cannot modify the FIPS mode if SNMPv3 users are already configured and present in the system. An error message is displayed if you attempt to change the FIPS mode by using the `fips mode enable` command in Global Configuration mode. You can enable or disable FIPS mode only if SNMPv3 users are not previously set up. If previously configured users exist on the system, you must delete the existing users before you change the FIPS mode.

Keep the following points in mind when you configure the AES128-CFB algorithm for SNMPv3:

1. SNMPv3 authentication provides only the `sha` option when FIPS mode is enabled.
2. SNMPv3 privacy provides only the `aes128 privacy` option when FIPS mode is enabled.
3. If you attempt to enable or disable FIPS mode and if any SNMPv3 users are previously configured, an error message is displayed stating you must delete all of the SNMP users before changing the FIPS mode.
4. A message is logged indicating whether FIPS mode is enabled for SNMPv3. This message is generated only when the first SNMPv3 user is configured because you can modify the FIPS mode only when users are not previously configured. This log message is provided to assist your system security auditing procedures.

snmp-server user (for AES128-CFB Encryption)

Specify that AES128-CFB encryption algorithm needs to be used for transmission of SNMP information. The Advanced Encryption Standard (AES) Cipher Feedback (CFB) 128-bit encryption algorithm is in compliance with RFC 3826. RFCs for SNMPv3 define two authentication hash algorithms, namely, HMAC-MD5-96 and HMAC-SHA1-96. These are the full forms or editions of the truncated versions, namely, HMAC-MD5 and HMAC-SHA1 authentication algorithms.

NOTE: Only the options that have been newly introduced are described here. For a complete description on all of the keywords and variables that are available with this command, refer the respective *Command Reference Guide* of the applicable platform of the Release 9.2.0.0 documentation set.

Z-Series S4810 S4820T S6000 MXL I/O Aggregator

Syntax

`snmp-server user` *`name`* `{`*`group_name`* `remote` *`ip-address`* `udp-port` *`port-number`*`} [1 | 2c | 3] [encrypted] [auth {md5 | sha}` *`auth-password`*`] [priv {des56 | aes128-cfb}` *`priv- password`*`] [access` *`access-list-name`* `| ipv6` *`access-list-name`* `|` *`access-list-name`* `ipv6` *`access-list-name`*`]`

To remove a user from the SNMP group, use the `no snmp-server user` *`name`* `{`*`group_name`* `remote` *`ip-address`* `udp-port` *`port-number`*`} [1 | 2c | 3] [encrypted] [auth {md5 | sha}` *`auth-password`*`] [priv {des56 | aes128-cfb}` *`priv-password`*`] [access` *`access-list-name`* `| ipv6` *`access-list-name`* `|` *`access-list-name`* `ipv6` *`access-list-name`*`]` command.

Parameters

auth-password (OPTIONAL) Enter a text string (up to 20 characters long) password that enables the agent to receive packets from the

	host and to send packets to the host. Minimum: eight characters long.
aes128	(OPTIONAL) Enter the keyword `aes128` to initiate the AES128-CFB encryption algorithm for transmission of SNMP packets.
priv-password	(OPTIONAL) Enter a text string (up to 20 characters long) password that enables the host to encrypt the contents of the message it sends to the agent and to decrypt the contents of the message it receives from the agent. Minimum: eight characters long.

Defaults

If no authentication or privacy option is configured, then the messages are exchanged (attempted anyway) without any authentication or encryption.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Added support for the AES128-CFB encryption algorithm on the S4820T, S4810, S6000, Z-Series, MXL, and I/O Aggregator platforms

Usage Information

To enable robust, effective protection and security for SNMP packets transferred between the server and the client, you can use the `snmp-server user` *`username`* `group` *`groupname`* `3 auth` *`authentication-type auth-password`* `priv aes128` *`priv-password`* to specify that AES128-CFB encryption algorithm needs to be used.

You cannot modify the FIPS mode if SNMPv3 users are already configured and present in the system. An error message is displayed if you attempt to change the FIPS mode by using the `fips mode enable` command in Global Configuration mode. You can enable or disable FIPS mode only if SNMPv3 users are not previously set up. Otherwise, you must remove the previously configured users before you change the FIPS mode.

Example

```
FTOS# snmp-server user privuser v3group v3 encrypted auth md5
9fc53d9d908118b2804fe80e3ba8763d priv aes128
d0452401a8c3ce42804fe80e3ba8763d
```

Related Commands

show snmp user — displays the information configured on each SNMP user name.

18

Stacking

This chapter describes the stacking enhancements and contains the following sections:

- Configuring the Uplink Speed of Interfaces as 40 Gigabit Ethernet

Configuring the Uplink Speed of Interfaces as 40 Gigabit Ethernet

You can configure the I/O Aggregator switch in standalone, VLT, and stack modes to operate with an uplink speed of 40 Gigabit Ethernet per second. Although the I/O Aggregator in programmable MUX mode supports the functionality to configure any base module port or optional Flex IO QSFP+ module port in native 40 GbE mode from Dell Networking OS Release 9.2.0.0, you can use the chassis management controller (CMC) interface to access the switch and specify the 40 GbE QSFP+ module ports to function in 40 GbE mode after the subsequent reload operation. By default, these QSFP+ modules function in 10GbE mode.

When you configure the native mode to be 40 GbE, the CMC sends a notification to the IOA to set the default internal working of all of the ports to be 40 GbE after the reload of the switch is performed. After you configure the native mode that denotes the uplink speed of the module ports to be 40 GbE, you must enter the `reboot` command (not pressing the Reset button, which causes the factory default settings to be applied when the device comes up online) from the CMC to cause the configuration of the uplink speed to be effective.

This functionality to set the uplink speed is available from the CLI or the CMC interface when the I/O Aggregator functions as a simple MUX or a VLT node with all of the uplink interfaces configured to be member links in the same LAG bundle. You cannot configure the uplink speed to be set as 40 GbE by default if the Aggregator functions in programmable MUX mode with mutiple uplink LAG interfaces or in stacking mode because CMC is not involved with configuration of parameters when the Aggregator operates in either of these modes with uplink interfaces being part of different LAG bundles.

After you restart the Aggregator, the 4-Port 10-Gigabit Ethernet modules or the 40GbE QSFP+ port that is split into four 10GbE SFP+ ports cannot be configured to be part of the same uplink LAG bundle that is set up with the uplink speed of 40 GbE. In such a condition, you can perform a hot-swap of the 4-port 10 Gbe Flex IO modules with a 2-port 40 GbE Flex IO module, which causes the module to become a part of the LAG bundle that is set up with 40 GbE as the uplink speed without another reboot. The Aggregator supports native 40 GbE mode for QSFP ports only in simple MUX mode and stacking mode of operation. In stacking mode, the base 40 GbE module ports are used for stacking and native 40 gbE uplink speed is enabled for only the QSFP+ ports on the optional 2-Port 40-Gigabit Ethernet QSFP+ FlexIO modules.

The following table describes the various speeds in different Aggregator modes. If a 4X10G SFP+ or a 4x10BASE-T module is plugged in and 40 Gbe mode is configured, it is in error-disabled state.

Table 9. Speeds in Different Aggregator Modes

Module Type	Standalone 10G mode	Standalone 40G Mode	Stacking 10G Mode	Stacking 40G mode	VLT 10G Mode	VLT 40G Mode
Base module	10G	40G	40G (HiGig)	40G	40G (Native)	40G

Module Type	Standalone 10G mode	Standalone 40G Mode	Stacking 10G Mode	Stacking 40G mode	VLT 10G Mode	VLT 40G Mode
				(HiGig)		(Native)
Optional module (2 40GbE)	10G	40G	10G	40G	10G	40G
Optional modules (4 10GbE)	10G	Error	10G	Error	10G	Error
FC module	10G	10G	10G	10G	10G	10G

To configure the uplink speed of the member interfaces in a LAG bundle for the Aggregator that operates in standalone, stacking, or VLT mode to be 40 Gigabit Ethernet per second, perform the following:

Specify the uplink speed of the member interfaces in a LAG bundle for the Aggregator that operates in standalone, stacking, or VLT mode to be 40 GbE. By default, the uplink speed of the LAG bundle is set as 10 GbE. You cannot configure the uplink speed if the Aggregator operates in programmable MUX mode. The `stack-unit unit-number iom-mode [stack | standalone | vlt]` 40G command is available the CMC interface and the CLI interface.

CONFIGURATION

`stack-unit` *`unit-number`* `iom-mode [stack | standalone | vlt] 40G`

You can use the `show system stack-unit` *`unit-number`* `iom-uplink-speed` command to view the uplink speed of the LAG bundles configured on the Flex IO modules installed on the Aggregator. The value under the Boot-speed field in the output of the show command indicates the uplink speed that is currently effective on the LAG bundles, whereas the value under the **Next-Boot** field indicates the uplink speed that is applicable for the LAG bundle after the next reboot of the switch. Depending on the uplink speed configured, the fan-out setting is designed accordingly during the booting of the switch.

The following example displays the output of the `show system stack-unit` *`unit-number`* `iom-uplink-speed` command with the Boot-speed field contained in it:

```
Dell# show system stack-unit 0 iom-uplink-speed
Unit    Boot-speed                      Next-Boot
-----------------------------------------------
0         10G                               40G
```

stack-unit iom-mode uplink-speed

Specify the uplink speed of the member interfaces in a LAG bundle for the Aggregator that operates in standalone, stacking, or VLT mode to be 40 GbE. By default, the uplink speed of the LAG bundle is set as 10 GbE.

Syntax

`stack-unit` *`unit-number`* `iom-mode {stack | standalone | vlt} uplink-speed 40G`

To restore the default uplink speed of the LAG bundle, which is 10 GbE, use the `stack-unit` *`unit-number`* `iom-mode {stack | standalone | vlt}` command.

Parameters

unit number <0-5>	Enter the number of the member stack unit. The range is from 0 to 5.

	iom-mode	Denotes the operating mode of the I/O Aggregator.
	stack	Specify that the uplink speed of the member interfaces in a LAG bundle applies for the Aggregator in stacking mode
	standalone	Specify that the uplink speed of the member interfaces in a LAG bundle applies for the Aggregator in standalone mode
	vlt	Specify that the uplink speed of the member interfaces in a LAG bundle applies for the Aggregator in VLT mode
	uplink-speed 40G	Set the uplink speed of the member or child interfaces of the LAG bundle to function at 40 Gigabit Ethernet per second
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator
Usage Information	This functionality to set the uplink speed is available from the CMC interface when the I/O Aggregator functions as a simple MUX or a VLT node with all of the uplink interfaces configured to be member links in the same LAG bundle. You cannot configure the uplink speed to be set as 40 GbE by default if the Aggregator functions in programmable MUX mode with mutiple uplink LAG interfaces or in stacking mode because CMC is not involved with configuration of parameters when the Aggregator operates in either of these modes with uplink interfaces being part of different LAG bundles. When you configure the native mode to be 40 GbE, the CMC sends a notification to the IOA to set the default internal working of all of the ports to be 40 GbE after the reload of the switch is performed. After you configure the native mode that denotes the uplink speed of the module ports to be 40 GbE, you must enter the `reboot` command (not pressing the Reset button, which causes the factory default settings to be applied when the device comes up online) from the CMC to cause the configuration of the uplink speed to be effective.	

show system stack-unit iom-uplink-speed

Display the uplink speed of the LAG bundles configured on the Flex IO modules installed on the Aggregator.

Syntax	`show system stack-unit` *`unit-number`* `iom-uplink-speed`	
Parameters	***unit number <0-5>***	Enter the number of the member stack unit. The range is from 0 to 5.
Command Modes	EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator

Usage Information	The value under the Boot-speed field in the output of the show command indicates the uplink speed that is currently effective on the LAG bundles, whereas the value under the Next-Boot field indicates the uplink speed that is applicable for the LAG bundle after the next reboot of the switch.

Example	`Dell# show system stack-unit 0 iom-uplink-speed`
	`Unit  Boot-speed  Next-Boot`
	`----------`
	`0  10G  40G`

stack-unit priority

Configure the ability of a switch to become the management unit of a stack.

Syntax `stack-unit` *`stack-number`* `priority` *`1-14`*

Parameters

stack-number	Enter the stack member unit identifier.
1–14	This preference parameter allows you to specify the management priority of one backup switch over another, with 0 the lowest priority and 14 the highest. The switch with the highest priority value is chosen to become the management unit if the active management unit fails or on the next reload.

Defaults 0

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator.

Related Commands

- reload – reboots FTOS.
- show system (S-Series) – displays the status of all stack members or a specific member.

stack-unit renumber

Change the stack member ID of any stack member or a stand-alone unit.

Syntax `stack-unit 0-11 renumber 0-11`

Parameters

0-11	The first instance of this value is the stack member unit identifier, from 0 to 11, of the switch that you want add to the stack. The range is: 0 to 11. The second instance of this value is the desired new unit identifier number.

Defaults none

Command Modes	EXEC Privilege
Command History	**Version 9.3.0.0** Introduced on the M I/O Aggregator
Usage Information	You can renumber any switch, including the management unit or a stand-alone unit. You cannot renumber a unit to a number of an active member in the stack. When executing this command on the master, the stack reloads. When the members are renumbered, only that specific unit is reset and comes up with the new unit number.

Example	`FTOS#stack-unit 5 renumber 6`
	
	`Renumbering will reset the unit.`
	`Warning: Interface configuration for current unit will be lost!`
	`Proceed to renumber [confirm yes/no]:`

Related Commands	• reload – reboots FTOS. • reset stack-unit – resets the designated S-Series stack member. • show system (S-Series) – displays the current status of all stack members or a specific member.

19

Virtual Link Trunking (VLT)

This chapter describes the VLT enhancements and contains the following sections:

- VLT Nodes as Rendezvous Points for Multicast Resiliency
- Specifying VLT Nodes in a PVLAN
- Proxy ARP Capability on VLT Peer Nodes

Specifying VLT Nodes in a PVLAN

You can configure VLT peer nodes in a private VLAN (PVLAN) on the S4810, S4820T, Z9000, and MXL platforms.

Virtual Link Trunking (VLT) is a mechanism that enables the physical links between two devices that are called VLT nodes or peers, and within a VLT domain, to be considered as a single logical link to external devices that are connected using LAG bundles to both the VLT peers. This capability enables redundancy without the implementation of Spanning Tree Protocol (STP), thereby providing a loop-free network with optimal bandwidth utilization.

You can configure the VLT peers or nodes in a private VLAN (PVLAN). Because the VLT LAG interfaces are terminated on two different nodes, PVLAN configuration of VLT VLANs and VLT LAGs are symmetrical and identical on both the VLT peers. PVLANs provide Layer 2 isolation between ports within the same VLAN. A PVLAN partitions a traditional VLAN into subdomains identified by a primary and secondary VLAN pair. With VLT being a Layer 2 redundancy mechanism, support for configuration of VLT nodes in a PVLAN enables Layer 2 security functionalities to be achieved. To enable maximum VLT resiliency to be obtained, you must configure the PVLAN IDs and mappings to be identical on both the VLT peer nodes.

The association of PVLAN with the VLT LAG must also be identical. After the VLT LAG is configured to be a member of a PVLAN so that it becomes member of either the primary or secondary PVLAN (which is associated with the primary), ICL becomes an automatic member of that PVLAN on both switches so that PVLAN data flow received on one VLT peer for a VLT LAG can also be transmitted on that VLT LAG from the peer.

You can associate either a VLT VLAN or a VLT LAG to a PVLAN. You must first configure the VLTi or a VLT LAG by using the `peer-link port-channel` *`id-number`* command or the VLT VLAN by using the `peer-link port-channel` *`id-number`* `peer-down-vlan vlan` *`interface number`* command and the `switchport` command. After you specify the VLTi link and VLT LAGs, you can associate the same port channel or LAG bundle that forms part of a VLT to a PVLAN by using the `interface` *`interface`* and `switchport mode private-vlan` commands.

When a VLT interconnect (VLTi) port in trunk mode is a member of symmetric VLT PVLANs using which PVLAN packets are traversed from one VLT node to the other, the PVLAN packets are forwarded only if the PVLAN settings of both the VLT nodes are identical. You can configure the VLTi in trunk mode to be a member of non-VLT PVLANs if the VLTi is configured on both the peers. MAC address synchronization is performed for VLT PVLANs across peers in a VLT domain.

Keep the following points in mind when you configure VLT nodes in a PVLAN:

- You must configure the VLTi link to be in trunk mode. You must not configure the VLTi link to be in access or promiscuous mode.

- You can configure a VLT LAG or port channel to be in trunk, access, or promiscuous port modes when you include the VLT LAG in a PVLAN. The VLT LAG settings must be the same on both the peers. If you configure a VLT LAG as a trunk port, you can associate that LAG to be a member of a normal VLAN or a PVLAN. If you configure a VLT LAG to be a promiscuous port, you can configure that LAG to be a member of PVLAN only. If you configure a VLT LAG to be in access port mode, you can add that LAG to be a member of secondary VLAN only.
- ARP entries are synchronized even when a mismatch occurs in the PVLAN mode of a VLT LAG.

Any VLAN that contains at least one VLT port as a member is treated as a VLT VLAN. You can configure a VLT VLAN to be a primary, secondary, or a normal VLAN. However, the VLT VLAN configuration must be symmetrical across peers. If the VLT LAG is tagged to any one of the secondary VLANs of a PVLAN or primary VLAN of a PVLAN, then both the primary and secondary VLANs are considered as VLT VLANs.

If you add an ICL or VLTi link as a member of a primary VLAN, the ICL becomes a part of the primary VLAN and its associated secondary VLANs, similar to the behavior for normal trunk ports. VLAN symmetricity is not validated if you associate an ICL to a PVLAN. Similarly, if you dissociate an ICL from a PVLAN, although the PVLAN symmetricity exists, ICL is removed from that PVLAN in such a case.

Association of VLTi as a Member of a PVLAN

If a VLAN is configured as a non-VLT VLAN on both the peers, the VLTi link is made as a member of that VLAN if the VLTi link is configured as a PVLAN/normal VLAN on both the peers. If a PVLAN is configured as a VLT VLAN on one peer and a non-VLT VLAN on another peer, the VLTi is added as a member of that VLAN by verifying the PVLAN symmetricity on both peers. In such a case, if a PVLAN is present as a VLT PVLAN on at least one of the peers, then symmetric configuration of PVLAN is validated to cause the VLTi to be a member of that VLAN. Whenever a change in the VLAN mode on one of the peers occurs, the information is synchronized with the other peer and VLTi is either added or removed from VLAN based on the validation of the LAN symmetricity.

For VLT VLANs, the association between primary VLAN and secondary VLAN is examined on both the peers. Only if the association is identical on both the peers, VLTi is configured as a member of those VLANs. This behavior is because of security functionalities in a PVLAN. If a VLAN is a primary VLT VLAN on one peer and not a primary VLT VLAN on the other peer, VLTi is not made a part of that VLAN. If a VLAN is secondary VLT VLAN on one peer and not a secondary VLT VLAN on the other peer, VLTi is not a part of that VLAN. If a VLAN is a normal VLT VLAN on one peer and a VLT PVLAN on the other peer, VLTi is not processed as a member of that VLAN.

MAC Synchronization for VLT Nodes in a PVLAN

For the MAC addresses that are learned on non-VLT ports, MAC address synchronization is performed with the other peer if the VLTi (ICL) link is part of the same VLAN as the non-VLT port. For MAC addresses that are learned on VLT ports, the VLT LAG mode of operation and the primary to secondary association of the VLT nodes is determined on both the VLT peers. MAC synchnronization is performed for the VLT LAGs only if the VLT LAG and primary-secondary VLT peer mapping are symmetrical.

The PVLAN mode of VLT LAGs on one peer is validated against the PVLAN mode of VLT LAGs on the other peer. MAC addresses are learned on that VLT LAG are synchronzied between the peers only if the PVLAN mode on both the peers is identical. If the MAC address is learned on a VLT LAG and the VLAN is a primary VLT VLAN on one peer and not a primary VLT VLAN on the other peer, MAC synchronization does not occur. Similarly, if the MAC address is learned on a VLT LAG and the VLAN is a secondary VLT VLAN on one peer and not a secondary VLT VLAN on the other peer, MAC synchronization does not occur. Additionally, if the MAc address is learned on a VLT LAG and the VLAN is a normal VLT VLAN on one peer and not a normal VLT VLAN on the other peer, MAC synchronization does not occur.

Whenever a change occurs in the VLAN mode of one of the peers, this modification in setting is synchronized with the other peers and depending on the validation mechanism that is initiated for MAC synchronization of VLT peers, MAC addresses learned on a particular VLAN are synchronized and made consistent to the other peers or MAC addresses synchronized from the other peers on the same VLAN are deleted. This method of processing occurs when the PVLAN mode of VLT LAGs is modified.

Because the VLTi link is only a member of symmetric VLT PVLANs, MAC synchronization takes place directly based on the membership of the VLTi link in a VLAN and the VLT LAG mode.

PVLAN Operations When One VLT Peer is Down

When a VLT port moves to the Admin or Operationally Down states on only one of the VLT nodes, the VLT Lag is still considered to be up. All the PVLAN MAC entries that correspond to the operationally down VLT LAG are maintained as synchronized entries in the device. These MAC entries are removed when the peer VLT LAG also becomes inactive or a change in PVLAN configuration occurs, which might cause inconsistency.

PVLAN Operations When a VLT Peer is Restarted

When the VLT peer node is rebooted, the VLAN membership of the VLTi link is preserved and when the peer node comes back online, a verification is performed with the newly received PVLAN configuration from the peer. If any differences are identified, the VLTi link is either added or removed from the VLAN. When the peer node restarts and returns online, all the PVLAN configurations are exchanged across the peers. Based on the information received from the peer, a bulk synchronization of MAc addresses that belong to spanned PVLANs is performed.

During the booting phase or when the ICL link attempts come up, a system logging message is recorded if VLT PVLAN mismatches, PVLAN mode mismatches, PVLAN association mismatches, or PVLAN port mode mismatches occur. Also, you can view these discrepancies if any that occur by using the `show vlt mismatch` command.

Interoperation of VLT Nodes in a PVLAN with ARP Requests

When an ARP request is received, the IP stack performs the following operations:

If the VLAN on which the ARP request is received is a secondary VLAN (community or isolated VLAN), if Layer 3 communication between secondary VLANs in a private VLAN is enabled by using the `ip local-proxy-arp` command in INTERFACE VLAN configuration mode, and if the ARP request is not received on the ICL, the ARP reply is sent with the MAC address of the primary VLAN. Additionally, an ARP request packet is originated on the primary VLAN for the intended destination IP address.

The ARP request received on ICLs are not proxied, even if they are received with a secondary VLAN tag, because the node from which the ARP request was forwarded would have replied with its MAC address and the current node discards the ARP request in such a case.

Scenarios for VLAN Membership and MAC Synchrnoization With VLT Nodes in PVLAN

The following table illustrates the manner in which association of VLTi link and PVLANs, and MAC synchronization of VLT nodes in a PVLAN is performed for various modes of operations of the VLT peers:

Table 10. VLAN Membership and MAC Synchrnoization With VLT Nodes in PVLAN

VLT LAG Mode		PVLAN Mode of VLT VLAN		ICL VLAN Membership	Mac Synchronization
Peer1	Peer2	Peer1	Peer2		
Trunk	Trunk	Primary	Primary	Yes	Yes
Trunk	Trunk	Primary	Normal	No	No
Trunk	Trunk	Normal	Normal	Yes	Yes
Promiscuous	Trunk	Primary	Primary	Yes	No
Trunk	Access	Primary	Secondary	No	No
Promiscuous	Promiscuous	Primary	Primary	Yes	Yes
Promiscuous	Access	Primary	Secondary	No	No
Promiscuous	Promiscuous	Primary	Primary	Yes	Yes
		- Secondary (Community)	- Secondary (Isolated)	No	No
Access	Access	Secondary (Community)	Secondary (Isolated)	No	No
		• Primary X	• Primary X	Yes	Yes
Promiscuous	Promiscuous	Primary	Primary	Yes	Yes
		- Secondary (Community)	- Secondary (Community)	Yes	Yes
		- Secondary (Isolated)	- Secondary (Isolated)	Yes	Yes
Promiscuous	Trunk	Primary	Normal	No	No
Promiscuous	Trunk	Primary	Primary	Yes	No
Access	Access	Secondary (Community)	Secondary (Community)	Yes	Yes
		- Primary VLAN X	- Primary VLAN X	Yes	Yes
Access	Access	Secondary (Isolated)	Secondary (Isolated)	Yes	Yes
		- Primary VLAN X	- Primary VLAN X	Yes	Yes
Access	Access	Secondary (Isolated)	Secondary (Isolated)	No	No
		- Primary VLAN X	- Primary VLAN Y	No	No
Access	Access	Secondary (Community)	Secondary (Community)	No	No
		- Primary VLAN Y	- Primary VLAN X	No	No
Promiscuous	Access	Primary	Secondary	No	No

VLT LAG Mode		PVLAN Mode of VLT VLAN		ICL VLAN Membership	Mac Synchronization
Peer1	Peer2	Peer1	Peer2		
Trunk	Access	Primary/Normal	Secondary	No	No

Configuring a VLT VLAN or LAG in a PVLAN

You can configure the VLT peers or nodes in a private VLAN (PVLAN). Because the VLT LAG interfaces are terminated on two different nodes, PVLAN configuration of VLT VLANs and VLT LAGs are symmetrical and identical on both the VLT peers. PVLANs provide Layer 2 isolation between ports within the same VLAN. A PVLAN partitions a traditional VLAN into subdomains identified by a primary and secondary VLAN pair. With VLT being a Layer 2 redundancy mechanism, support for configuration of VLT nodes in a PVLAN enables Layer 2 security functionalities to be achieved. This section contains the following topics that describe how to configure a VLT VLAN or a VLT LAG (VLTi link) and assign that VLT interface to a PVLAN.

Creating a VLT LAG or a VLT VLAN

1. Configure the port channel for the VLT interconnect on a VLT switch and enter interface configuration mode.

   CONFIGURATION mode

   `interface port-channel` *`id-number`*

   Enter the same port-channel number configured with the `peer-link port-channel` command as described in Enabling VLT and Creating a VLT Domain.

   NOTE: To be included in the VLTi, the port channel must be in Default mode (`no switchport` or VLAN assigned).

2. Remove an IP address from the interface.

   INTERFACE PORT-CHANNEL mode

   `no ip address`

3. Add one or more port interfaces to the port channel.

   INTERFACE PORT-CHANNEL mode

   `channel-member` *`interface`*

   `interface`: specify one of the following interface types:

   - 1-Gigabit Ethernet: Enter `gigabitethernet` *`slot/port`*.
   - 10-Gigabit Ethernet: Enter `tengigabitethernet` *`slot/port`*.

4. Ensure that the port channel is active.

   INTERFACE PORT-CHANNEL mode

   `no shutdown`

5. Repeat Steps 1 to 4 on the VLT peer switch to configure the VLT interconnect.

6. Enter VLT-domain configuration mode for a specified VLT domain.

   CONFIGURATION mode

   `vlt domain` *`domain-id`*

   The range of domain IDs is from 1 to 1000.

7. Enter the port-channel number that acts as the interconnect trunk.
   VLT DOMAIN CONFIGURATION mode
   `peer-link port-channel id-number`
   The range is from 1 to 128.
8. (Optional) To configure a VLT LAG Enter the VLAN ID number of the VLAN where the VLT forwards packets received on the VLTi from an adjacent peer that is down.
   VLT DOMAIN CONFIGURATION mode
   `peer-link port-channel id-number peer-down-vlan vlan interface number`
   The range is from 1 to 4094.

Associating the VLT LAG or VLT VLAN in a PVLAN

1. Access INTERFACE mode for the port that you want to assign to a PVLAN.
   CONFIGURATION mode
   `interface interface`
2. Enable the port.
   INTERFACE mode
   `no shutdown`
3. Set the port in Layer 2 mode.
   INTERFACE mode
   `switchport`
4. Select the PVLAN mode.
   INTERFACE mode
   `switchport mode private-vlan {host | promiscuous | trunk}`
   - `host` (isolated or community VLAN port)
   - `promiscuous` (intra-VLAN communication port)
   - `trunk` (inter-switch PVLAN hub port)
5. Access INTERFACE VLAN mode for the VLAN to which you want to assign the PVLAN interfaces.
   CONFIGURATION mode
   `interface vlan vlan-id`
6. Enable the VLAN.
   INTERFACE VLAN mode
   `no shutdown`
7. To enable maximum VLT resiliency to be obtained, you must configure the PVLAN IDs and mappings to be identical on both the VLT peer nodes. Set the PVLAN mode of the selected VLAN to primary.
   INTERFACE VLAN mode
   `private-vlan mode primary`

8. Map secondary VLANs to the selected primary VLAN.

   INTERFACE VLAN mode

   `private-vlan mapping secondary-vlan` *vlan-list*

   The list of secondary VLANs can be:

   – Specified in comma-delimited (*VLAN-ID,VLAN-ID*) or hyphenated-range format (*VLAN-ID-VLAN-ID*).
   – Specified with this command even before they have been created.
   – Amended by specifying the new secondary VLAN to be added to the list.

show vlt private-vlan

Display the association of private VLAN (PVLAN) with the VLT LAG. You can configure VLT peer nodes in a PVLAN on the S4810, S4820T, Z9000, and MXL platforms.

Syntax	`show vlt private-vlan`	
Command Modes	EXEC	
Command History	**Version 9.3.0.0**	Introduced on the Z9000, S4810, S4820T, and MXL platforms.
Usage Information	If you add an ICL or VLTi link as a member of a primary VLAN, the ICL becomes a part of the primary VLAN and its associated secondary VLANs, similar to the behavior for normal trunk ports. VLAN symmetricity is not validated if you associate an ICL to a PVLAN. Similarly, if you dissociate an ICL from a PVLAN, although the PVLAN symmetrictiy exists, ICL is removed from that PVLAN in such a case. The **ICL Status** field denotes the type of the VLAN port of the VLTi link configured in a PVLAN.	

Example

```
FTOS#Show vlt private-vlan vlan-id

Codes: C- Community, I - Isolated, V - Internally tagged, T -
tagged, * - VLT Pvlan
Primary    Secondary    ICL Status
10                        V (*)
            20(C)         V
            30 (I)       V

40                          T
            50(C)         T
            60 (I)       T


FTOS#
```

Proxy ARP Capability on VLT Peer Nodes

The proxy ARP functionality on VLT peer nodes is supported on the S4810, S4820T, Z9000, I/O Aggregator, and MXL platforms.

Proxy ARP enables hosts with knowledge of the network to accept and forward packets from hosts that contain no knowledge of the network. Proxy ARP makes it possible for hosts to be ignorant of the network, including subnetting.

Virtual Link Trunking (VLT) is a mechanism that enables the physical links between two devices that are called VLT nodes or peers, and within a VLT domain, to be considered as a single logical link to external devices that are connected using LAG bundles to both the VLT peers. This capability enables redundancy without the implementation of Spanning Tree Protocol (STP), thereby providing a loop-free network with optimal bandwidth utilization.

A Proxy ARP-enabled device answers the ARP requests that are destined for another host or router. This phenomenon operates by causing the local host to consider that the Proxy ARP-enabled device is the originator or the owner of the IP address, and the local host forwards the traffic to the proxy ARP-enabled device, which in turn transmits the packets to the real destination.

By default, proxy ARP is enabled. To disable Proxy ARP, use the `no proxy-arp` command in the interface mode. To re-enable Proxy ARP, use the `ip proxy-arp` command in INTERFACE mode. To view if Proxy ARP is enabled on the interface, use the `show config` command in INTERFACE mode. If it is not listed in the `show config` command output, it is enabled. Only non-default information is displayed in the show config command output.

ARP proxy operation is performed for the IP address of the peer VLT node when the peer VLT node is down. The working of ARP proxy is stopped either when the peer routing timer expires or when the peer VLT node goes up. Layer 3 VLT provides a higher resiliency at the Layer 3 forwarding level. VLT peer routing enables you to replace VRRP with routed VLT to route the traffic from Layer 2 access nodes. With proxy ARP, hosts can resolve the MAc address of the VLT node even when VLT node is down.

When a VLT node receives an ARP request for the IP address of the VLT peer, owing to LAG-level hashing algorithm in the top-of-rack (TOR) switch, the incorrect VLT node responds to the ARP request with the peer MAC address if the ICL link is down. Proxy ARP is not performed when the ICL link is up and the ARP request the wrong VLT peer. In this case, ARP requests are tunneled to the VLT peer.

Proxy ARP supported on both VLT interfaces and non-VLT interfaces. Proxy ARP supported on symmetric VLANs only. Proxy ARP is enabled by default. Routing table must be symmetrically configured to support proxy ARP. For example, connsider a sample topology in which VLAN 100 is configured on two VLT nodes, node 1 and node 2. ICL link is not configured between the two VLT nodes. Assume that the IP address of VLAN 100 in node 1 is 10.1.1.1/24 and IP address of VLAN 100 in node 2 is 20.1.1.2/24. In this case i,f the ARP request for 20.1.1.1 reaches node 1, Node 1 will not perform ARP for 20.1.1.2. Proxy ARP is supported only for the IP address belongs to the received interface IP network. Proxy ARP is not supported if the ARP requested IP address is different from the received interface IP subnet. For example, if VLAN 100 and 200 is configured on the VLT peers, and if IP address of VLAN 100 is configured as 10.1.1.0/24 and IP address of VLAN 200 is configured as 20.1.1.0/24, Proxy ARP is not performed if the VLT node receives ARP request for 20.1.1.0/24 on VLAN 100.

Working of Proxy ARP for VLT Peer Nodes

Proxy ARP is enabled only when peer routing is enabled on both the VLT peers. If peer routing is disabled on one of the VLT peers, proxy ARP is not performed when ICL link goes down. Proxy ARP is performed only when the VLT peer's MAC address is installed in the database. Proxy ARP is stopped when the VLT peer's MAC address is removed from the ARP database owing to peer routing timer expiry. The source

hardware address in the ARP response contains the VLT peer MAC address. Proxy ARP is supported for both unicast and broadcast ARP requests. Control packets other than ARP requests destined to the VLT peers that reach the undesired and incorrect VLT node are dropped if the ICL link is down. Further processing is not done on these control packets. VLT node does not perform any action if it receives gratuitous ARP requests for the VLT peer IP address. Proxy ARP is also supported on secondary VLANs. If VLT nodes are configured with private VLANs, and the ARP request for private VLAN IP address reaches the wrong peer, when the ICL link or peer is down, then the wrong peer responds to the ARP request with the peer MAC address.

IP address of the VLT node VLAN interfaces is synchronized with the VLT peer over ICL when VLT peers are up. Whenever an IP address is added or deleted, this updated information is synchronized with the VLT peer. IP address synchronization occurs regardless of the VLAN administrative state. IP address addition and deletion serve as the trigger events for synchronization. When a VLAN state is down, the VLT peer might perform proxy ARP operation for the IP addresses of that VLAN interface.

VLT nodes start performing Proxy ARP when the ICL link goes down. When the VLT peer becomes operationally up, proxy ARP will be stopped for the peer VLT IP addresses. When the peer node is rebooted, the IP address synchronized with the peer is not flushed. Peer down events cause proxy ARP to be commenced.

When a VLT node detects peer up, it will not perform proxy ARP for the peer IP addresses. IP address synchronization occurs again between the VLT peers.

Proxy ARP is enabled only if peer routing is enabled on both the VLT peers. If you disable peer routing by using the `no peer-routing` command in VLT DOMAIN node, a notification is sent to the VLT peer to disable the proxy ARP. If peer routing is disabled when ICL link is down, a notification is not sent to the VLT peer and in such a case, the VLT peer does not disable the proxy ARP operation.

When VLT domain is removed on one of the VLT nodes, the peer routing configuration removal will be notified to the peer. In this case VLT peer node will disable the proxy ARP. When the ICL link is removed on one of the VLT nodes by using the `no peer-link` command, the ICL down event is triggered on the other VLT node, which in turn starts the proxy ARP application. The VLT node where the ICL link is deleted flushes the peer IP addresses and does not perform proxy ARP for the further LAG hashed ARP requests.

VLT Nodes as Rendezvous Points for Multicast Resiliency

You can configure virtual link trunking (VLT) peer nodes as rendezvous points (RPs) in a Protocol Independent Multicast (PIM) domain on the S4810, S4820T, and Z9000 platforms. This capability enables VLT resiliency and robustness for multicast routing operations.

PIM uses a VLT node as the root or a Rendezvous Point (RP) of the share tree distribution tree to distribute multicast traffic to a multicast group. Messages to join the multicast group (Join messages) are sent towards the RP and data is sent from senders to the RP so receivers can discover who are the senders and begin receiving traffic destined to the multicast group.

To enable an explicit multicast routing table synchronization method for VLT nodes, you can configure VLT nodes as RPs. Because multicast routing requires the incoming interface for each route to be identified, PIM running on both VLT peers enables both the peers to obtain traffic from the same incoming interface.

You can configure a VLT node to be an RP by entering the `ip pim rp-address` command in Global Configuration mode. When you configure a VLT node as an RP, the (*, G) routes that are synchronized from the VLT peers are ignored and not downloaded to the device. For the (S, G) routes that are synchronized from the VLT peer, after the RP starts receiving multicast traffic via the (S,G), these (S, G) routes are considered valid and are downloaded to the device. Only (S,G) routes are used to forward the multicast traffic from the source to the receiver.

You can configure VLT nodes that function as RP as Multicast Source Discovery Protocol (MSDP) peers in different domains. However, you cannot configure the VLT peers as MSDP peers in the same VLT domain. In such a case, RP functionality is not supported by the VLT peer.

If the same source or RP is can be accessed over both a VLT and a non-VLT VLAN, you must configure better metrics for the VLT VLANs. Otherwise, it is possible that one VLT node chooses a non-VLT VLAN (if the path through the VLT Vlan was not available when the route was learnt) and another VLT node selects a VLT Vlan. Such a scenario can cause duplication of packets. ECMP is not supported when you configure VLT nodes as RPs.

Backup RP is not supported if VLT peer that functions as the RP is statically configured. With static RP configuration, if the RP reboots, it can handle new clients only after it comes back online. Until the RP returns to the active state, the VLT peer forwards the packets for the already logged-in clients. To enable the VLT peer node to retain the synchronized multicast routes or synchronized multicast outgoing interface (OIF) maps after a peer node failure , the timeout value that you configured by using the `multicast peer-routing timeout` *`value`* command is used. You can configure the time for a VLT node to retain synced multicast routes or synced multicast outgoing interface (OIF) after a VLT peer node failure by using the `multicast peer-routing-timeout` command in VLT DOMAIN mode to be an optimal value. Using the bootstrap router (BSR) mechanism, both the VLT nodes in a VLT domain can be configured as the candidate RP for the same group range. When an RP fails, the VLT peer automatically takes over the role of the RP. This phenomenon enables resiliency to be achieved by the PIM BSR protocol.

20

Documentation Updates

Because the entire hardware and software documentation set is not being published for Release 9.3.0.0, similar to the earlier major releases, this section has been organized to present behavioral-changes and enhancements to commands and configuration settings that have been implemented in Release 9.3.0.0. We recommend that you read this section in conjunction with the full-blown documentation set of Release 9.2.0.0 of the different platforms.

- Starting with Dell Networking OS Release 9.3.0.0, on an I/O Aggregator, you can enable and disable the DCB capability without having to reload the switch. The `on-next-reload` option available with the `[no] dcb enable` command is not required when you enter this command. Until Release 9.2.0.0, you could enable or disable DCB only by specifying the `on-next-reload` keyword, which caused DCB to be activated or deactivated only during the subsequent reload of the switch. Similarly, you can configure automatic detection of DCB and DCBx by using the `dcb enable auto-detect` command without the `on-next-reload` keyword.

  The enhanced syntax of the `dcb enable` command is as follows:

  ```
  [no] dcb enable [auto-detect]
  ```

  The enhanced syntax of the dcb enable command is as follows:

  ```
  dcb enable auto-detect [on-next-reload]
  ```

- The following note applies to the *Enhanced VLT section of the VLT* chapter of the relevant platforms' Configuration Guides:

  In an enhanced VLT (eVLT) configuration, if different virtual router identifiers (VRIDs) are configured in the two VLT domains that are part of the same Layer 2 domain, and if you enable generation of SNMP traps for VRRP in the devices of both the VLT domains by using the `snmp-server enable traps vrrp` command, a large number of traps are received on the device that receives the VRRP traps for state changes because traps are generated every second and the alarm trap application is throttled.

- The following note applies to the `rate-shape` command topic of the QoS chapter of the S6000 Command Reference Guide:

  On the S6000 platform, hierarchical QoS does not support policy-based rate shaping if you configure a unicast queue as a strict-priority (SP) queue.

- The following note applies to the Private VLAN Concepts section of the PVLAN chapter of all platforms' Configuration Guides:

  On the S4810 platform, stale and incorrect entries are added to the MAC address table and CAM when movement of multiple stations occurs across the primary and secondary private VLANs. Because PVLAN uses VLAN translation tables, the station movement is not detected if such a shift of station occurs in a single PVLAN domain.

- The following additional information applies to the Private VLAN Concepts section of the PVLAN chapter of all platforms' Configuration Guides:

  In a VLAN domain that contains two VLT peer nodes connected by the VLT interconnect link (VLTi) and with IGMP snooping and PIM enabled in the VLAN, IGMP queries generated by one peer are sent to the other peer that contains an orphan port and forwaded out of it. This behavior occurs when the peer that functions as the IGMP querier contains an extended ACL with deny action configured and

applied on the egress interface, which must not cause the IGMP queries to be transmitted out of the interface.

- The following note applies to the Configuring a VLT VLAN or LAG in a PVLAN section of this Addendum document:

  When you configure VLT nodes in a PVLAN, you must create the PVLAN after you configure a VLT domain. If you configure the PVLAN before you set up the VLT domain, you must delete and reconfigure the VLANs to enable the VLT interconnect (VLTi) or ICL membership to be correctly added. Otherwise, the ICL is not properly tagged with the PVLAN.

- The following note applies to the Microsoft Network Load Balancing overview topic of this Addendum document:

  The flow processor (FP) entry that is used for Microsoft NLB is not redefined after a reload of the system if the interface that is used for NLB was shut down before the system reload. This problem occurs because the `arp ip-address multicast-mac-address` command that you configure to associate an IP address with a multicast MAC address retrieves an interface port and slot from the interface list specified using the Layer 2 multicast command of `mac-address-table static` command. As a result, the interface configured in static ARP is mapped with that of the L2 multicast configuration and matched with the associated VLAN ID.

  When you shut down an interface and reload a system, all ARP entries are reset and the interface FP entry is removed. After the system comes up, no IFP entry is associated with the static NLB ARP. This behavior occurs because the interface is in shutdown state and hence no active ARP entry is maintained by the ARP application. Therefore, you must always associate the NLB configuration with one of the active or enabled interfaces.

- The following note applies to the Priority-Based Flow Control section of the DCB chapter of I/O Aggregator Configuration Guide:

  On an I/O Aggregator, PFC is enabled by default with a priority value of 4 if you enable automatic detection of DCB with PFC turned off. This behavior occurs if you perform a reload of the switch with the `dcb enable auto-detect on-next-reload` command configured. The output of the `show interface interface-type slot/port pfc details` command displays the priority list as 4.

- The following note is additional information to the Configuring PFC section of the DCB chapter of all of the platforms' Configuration Guides:

  When you apply or remove a DCB input policy from an interface, one or two CRC errors are expected to be noticed on the ingress ports for each removal or attachment of the policy. This behavior occurs because the port is brought down when PFC is configured. When a DCB input policy with PFC profile is configured or unconfigured on an interface or a range of interfaces not receiving any traffic, interfaces with PFC settings and that receive appropriate PFC-enabled traffic (unicast, mixed-frame-size traffic) display incremented values in the "CRC" and "Discards" counters. (These ingress interfaces receiving pfc-enabled traffic have an egress interface that has a compatible PFC configuration).

- The following note applies to the CAM Profiles section of the Content Addressable Memory (CAM) chapter of the MXL, S4810/4820T, Z9000 configuration guides:

  If you define a CAM profile in which the CAM block size for IPv4 is specified as zero on a VLT peer, duplicate multicast data packets are obtained by the receivers. This condition occurs owing to a missing egress mask entry that is not installed in the hardware table because of the IPv4 table size being zero.

- The following note applies to the Private VLAN Concepts section of the PVLAN chapter of all platforms' Configuration Guides:

  PVLAN uses VLAN translation tables. This mechanism does not enable movement of stations to be detected if the shift of the station occurs in a single PVLAN domain. It is detected correctly only with a top-of-rack (ToR) node.

- The following note applies to the Marking Egress Packets with a DEI Value section of the S4810, S4820T, S6000, Z9000, and MXL Switch Configuration Guides:

  You cannot set the Canonical Format Identifier (CFI) bit alone in an outgoing packet because it will cause the IEEE 802.1ad drop eligible indicator (DEI) bit in the outgoing dot1p packet to also be reset. This behavior is expected when you configure the incoming DEI value to be honored by mapping it to an FTOS drop precedence, and also specify the DEI marking on the ingress and egress interface. For example, if you enter the `enable dei honor 0` yellow command on the ingress interface and the `enable dei mark yellow 1` on both the ingress and egress interfaces, and send traffic with CFI bit set to 0, the CFI bit is set to 0 instead of being set to 1 on the egress interface.

- The following note applies to the Configure Layer 2 and Layer 3 ACLs section of the Access Control Lists (ACLs) chapter of the Configuration Guides of all of the platforms:

  The egress MAC ACL counter shows zero for existing rules in a VLAN, when a new rule is added with a sequence number that is less than the existing rule. The counter of the old rule is reset to zero. The counter of the old rule is updated starting with zero, even though the lesser sequence number rule is removed from the access group. The update time of the counter differs for continuous and discrete traffic for VLAN, physical, and port channel interfaces. This behavior is seen with both standard and extended MAC ACLs on the ingress and egress direction.

- The following additional information applies to the Important Points to Remember section of the Virtual Link Trunking (VLT) chapter of the S4820T Configuration Guide:

  Some of the IPv6 neighbors and ARP entries are learned on the VLT LAG member (physical port) and VLTi instead of being learned on the VLT LAG after reloading the TOR device. The IP application receives ARP and neighbor advertisement (NA) packets over the physical port. It learns ARP/NA entries over the physical port.

- The following information regarding the enabling of DCB globally on a system applies to the `dcb-enable` command topic of the Data Center bridging (DCB) chapter of the S6000 Command Reference Guide:

  On an S6000 Switch, when BFD or LACP with faster convergence is enabled and if you enter the `dcb-enable` command to enable DCB globally, a protocol flap is an expected behavior.

- The following note applies to the `rate-shape` command topic of the QoS chapter of the S6000 Command Reference Guide:

  When you configure per-queue level rate-shaping, do not use values less than 100 Mbps.

- The following note applies to the Proxy ARP section of the VLT chapter of this addendum document:

  Because the proxy ARP capability depends on the peer routing settings, you must ensure the following: You must remove the VLT domain in both the peers. You must configure the VLT peer routing timeout by using the `multicast peer-routing-timeout` command in VLT DOMAIN mode as a definitive value instead of an infinite value.

- The following note applies to the `service-class dot1p-mapping` command topic of the QoS chapter of the MXL Command Reference Guide:

  With default dot1p to queue mapping on an MXL Switch, ETS configuration from Cisco Nexus 7000 Series Switches not supported. In such an environment, you must add the `service-class dot1p-mapping dot1p2 2` command to the configuration file of the MXL Switch to enable the default Cisco FCoE service policy to work when service-class dynamic dot1p is used.

- The following note applies to the Guidelines section of the FC Flex IO chapter of this addendum document:

  While storage traffic traverses through FC Flex IO modules and the Ethernet uplink port-channel status changes (with DCB enabled on an adjacent switch), FCoE traffic is disrupted. This disruption

does not occur if there is no Ethernet traffic and only FCoE traffic is present, or if DCB remains disabled on the ToR switch.

- The following note applies to the Protocol Separation section of the EIS chapter of the addendum document:

  If you configure a source interface is for any EIS management application, EIS might not coexist with that interface and the behavior is undefined in such a case. You can configure the source interface for the following applications: FTP, ICMP (ping and traceroute utilites), NTP, RADIUS, TACACS, Telnet, TFTP, syslog, and SNMP traps. Out of these applications, EIS can coexist with syslog and SNMP traps only because these applications do not require a response after a packet is sent.

- The following note applies to the Config Storm Control section of Storm Control chapter of the S6000 Configuration Guide:

  When both rate policing and storm control are configured, packets are dropped and are not sent out of the egress interface. Rate policing and storm control are mutually exclusive operations; you can configure only one of the two functionalities at a point in time. You can configure storm control at interface level or globally and it is restricted against the presence of the per-interface-level rate-police or the policy-level rate-police setting.

- The following note applies to to the VLT Nodes in PVLANs section of the addendum document: ARP entries are synchronized even when a mismatch occurs in the PVLAN mode of a VLT LAG. PR 128799

- The following note applies to the Post Configuration Script - BMP Mode section of the Configuring BMP chapter of the Open Automation Guide:

  Because the rstimer utility supports only the minute-level precision, the rstimer utility cannot be used for the first time after the ninth minute of the execution of the script. As a result, you cannot extend the preconfig or postconfig timer during the last minute of expiration of the timer.

- The following note applies to the Specifying an Auto-Failover Limit section of the High Availability (HA) chapter of the S4810 and S4820T Configuration Guides:

  Recovering the Switch When Auto-Failover Limit is Exceeded

  When the auto-failover limit is exceeded, the switches become disabled. In such a state, you must reboot the switches to be able to access the device again because you cannot use the preconfigured user credentials to log in to the device until it is reset. You must use the `reload` command to reboot the switches, including the management unit and the stack unit when the auto-failover limit is exceeded.

Configuring the Commands Without a Separate User Account for the PMUX Mode of the I/O Aggregator

Starting with Dell Networking OS Release 9.3.0.0, you can configure and specify the commands that were available in the programmable MUX mode of the I/O Aggregator until Release 9.2.0.0 without having to configure a user profile to access the PMUX mode. As a result, you do not need to define separate user accounts with permissions to access the PMUX mode on the switch. The user profile that you defined to access and log in to the switch is sufficient to configure these commands.

This part contains chapters that describe the commands that were previously available in the PMUX mode (until Release 9.2.0.0), which you can now configure without the need for an exclusive, separate user account. This chapter describes the commands and configuration settings that you can employ from the PMUX mode of the I/O Aggregator. Although these commands are indicated as having been introduced in Release 9.2(0.0), you had to specify a separate user profile with permissions to access PMUX mode to be able to use these commands until Release 9.2(0.0). Starting with Release 9.3(0.0), you can operate these commands and attributes without the need for a separate user profile.

21

Data Center Bridging (DCB)

Data center bridging (DCB) refers to a set of IEEE Ethernet enhancements that provide data centers with a single, robust, converged network to support multiple traffic types, including local area network (LAN), server, and storage traffic.

The Dell Networking operating software (Dell) commands for data center bridging features include 802.1Qbb priority-based flow control (PFC), 802.1Qaz enhanced transmission selection (ETS), and the data center bridging exchange (DCBX) protocol.

advertise dcbx-appln-tlv

On a DCBX port with a manual role, configure the application priority TLVs advertised on the interface to DCBX peers.

Syntax

`advertise dcbx-appln-tlv {fcoe | iscsi}`

To remove the application priority TLVs, use the `no advertise dcbx-appln-tlv {fcoe | iscsi}` command.

Parameters

{fcoe | iscsi} Enter the application priority TLVs, where:

- `fcoe`: enables the advertisement of FCoE in application priority TLVs.
- `iscsi`: enables the advertisement of iSCSI in application priority TLVs.

Defaults

Application priority TLVs are enabled to advertise FCoE and iSCSI.

Command Modes

PROTOCOL LLDP

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To disable TLV transmission, use the `no` form of the command; for example, `no advertise dcbx-appln-tlv iscsi`.

advertise dcbx-tlv

On a DCBX port with a manual role, configure the PFC and ETS TLVs advertised to DCBX peers.

Syntax

`advertise dcbx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]`

To remove the advertised ETS TLVs, use the `no advertise dcbx-tlv` command.

Parameters	**{ets-conf | ets-reco | pfc}**	Enter the PFC and ETS TLVs advertised, where: • `ets-conf`: enables the advertisement of ETS configuration TLVs. • `ets-reco`: enables the advertisement of ETS recommend TLVs. • `pfc`: enables the advertisement of PFC TLVs.
Defaults	All PFC and ETS TLVs are advertised.	
Command Modes	PROTOCOL LLDP	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	You can configure the transmission of more than one TLV type at a time; for example: `advertise dcbx-tlv ets-conf ets-reco`. You can enable ETS recommend TLVs (`ets-reco`) only if you enable ETS configuration TLVs (`ets-conf`). To disable TLV transmission, use the `no` form of the command; for example, `no advertise dcbx-tlv pfc ets-reco`. DCBX requires that you enable LLDP to advertise DCBX TLVs to peers. Configure DCBX operation at the INTERFACE level on a switch or globally on the switch. To verify the DCBX configuration on a port, use the `show interface dcbx detail` command.	

bandwidth-percentage

Configure the bandwidth percentage allocated to priority traffic in port queues.

Syntax	`bandwidth-percentage` *`percentage`* To remove the configured bandwidth percentage, use the `no bandwidth-percentage` command.	
Parameters	***percentage***	(Optional) Enter the bandwidth percentage. The percentage range is from 1 to 100% in units of 1%.
Defaults	none	
Command Modes	QOS-POLICY-OUT-ETS	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information	By default, equal bandwidth is assigned to each port queue and each dot1p priority in a priority group. To configure bandwidth amounts in associated dot1p queues, use the `bandwidth-percentage` command. When specified bandwidth is assigned to some port queues and not to others, the remaining bandwidth (100% minus assigned bandwidth amount) is equally distributed to unassigned nonstrict priority queues in the priority group. The sum of the allocated bandwidth to all queues in a priority group must be 100% of the bandwidth on the link. ETS-assigned bandwidth allocation applies only to data queues, not to control queues. The configuration of bandwidth allocation and strict-queue scheduling is not supported at the same time for a priority group. If you configure both, the configured bandwidth allocation is ignored for priority-group traffic when you apply the output policy on an interface. By default, equal bandwidth is assigned to each priority group in the ETS output policy applied to an egress port if you did not configure bandwidth allocation. The sum of configured bandwidth allocation to dot1p priority traffic in all ETS priority groups must be 100%. Allocate at least 1% of the total bandwidth to each priority group and queue. If bandwidth is assigned to some priority groups but not to others, the remaining bandwidth (100% minus assigned bandwidth amount) is equally distributed to nonstrict-priority groups which have no configured scheduler.
Related Commands	• qos-policy-output ets — creates a QoS output policy. • scheduler — schedules priority traffic in port queues.

dcb-enable

Enable data center bridging.

Syntax	`dcb enable` To disable DCB, use the `no dcb enable` command.
Defaults	none
Command Modes	CONFIGURATION
Command History	**Version 9.2(0.0)** Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only. **Version 8.3.16.1** Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	DCB is not supported if you enable `link-level flow control` on one or more interfaces.

dcb-input

To apply pause or flow control for specified priorities using a configure delay time, create a DCB input policy.

Syntax

`dcb-input policy-name`

To delete the DCB input policy, use the `no dcb-input` command.

Parameters

policy-name	Maximum: 32 alphanumeric characters.

Defaults

none

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

As soon as you apply a DCB policy with PFC enabled on an interface, DCBx starts exchanging information with PFC-enabled peers. The IEEE802.1Qbb, CEE, and CIN versions of PFC TLV are supported. DCBx also validates PFC configurations received in TLVs from peer devices.

By applying a DCB input policy with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, also enable PFC on all DCB egress ports or configure the dot1p priority-queue assignment of PFC priorities to lossless queues (refer to `pfc no-drop queues`).

To remove a DCB input policy, including the PFC configuration it contains, enter the `no dcb-input policy-name` command in Interface Configuration mode.

Related Commands

dcb-policy input — applies the input policy with the PFC configuration.

dcb-output

To associate an ETS configuration with priority traffic, create a DCB output policy.

Syntax

`dcb-output policy-name`

To remove the ETS output policy globally, use the `no dcb output policy-name` command.

Parameters

policy-name	Enter the DCB output policy name. The maximum is 32 alphanumeric characters.

Defaults

none

Command Modes	CONFIGURATION	
Command History	**Version 8.3.16.1**	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	To associate a priority group with an ETS output policy with scheduling and bandwidth configuration, create a DCB output policy. You can apply a DCB output policy on multiple egress ports. When you apply an ETS output policy on an interface, ETS-configured scheduling and bandwidth allocation take precedence over any configured settings in QoS output policies. The ETS configuration associated with 802.1 priority traffic in a DCB output policy is used in DCBX negotiation with ETS peers.	
Related Commands	dcb-policy output — applies the output policy.	

dcb-policy input

Apply the input policy with the PFC configuration to an ingress interface.

Syntax	`dcb-policy input` *policy-name* To delete the input policy, use the `no dcb-policy input` command.	
Parameters	***policy-name***	Enter the input policy name with the PFC configuration to an ingress interface.
Defaults	none	
Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	If you apply an input policy with PFC disabled (`no pfc mode on`): • You can enable link-level flow control on the interface. To delete the input policy, first disable link-level flow control. PFC is then automatically enabled on the interface because an interface is by default PFC-enabled. • PFC still allows you to configure lossless queues on a port to ensure no-drop handling of lossless traffic. When you apply an input policy to an interface, an error message is displayed if: • The PFC dot1p priorities result in more than two lossless port queues globally on the switch. • You already enabled link-level flow control. PFC and link-level flow control cannot be enabled at the same time on an interface. In a switch stack, configure all stacked ports with the same PFC configuration.	

A DCB input policy for PFC applied to an interface may become invalid if you reconfigure the dot1p-queue mapping. This situation occurs when the new dot1p-queue assignment exceeds the maximum number (2) of lossless queues supported globally on the switch. In this case, all PFC configurations received from PFC-enabled peers are removed and resynchronized with the peer devices.

Traffic may be interrupted when you reconfigure PFC no-drop priorities in an input policy or reapply the policy to an interface.

Related Commands dcb-input — creates a DCB input policy.

dcb-policy input stack-unit stack-ports all

Apply the specified DCB input policy on all ports of the switch stack or a single stacked switch.

Syntax

```
dcb-policy input stack-unit {all | stack-unit-id} stack-ports
all dcb-input-policy-name
```

To remove all DCB input policies applied to the stacked ports and rest the PFC to its default settings, use the `no dcb-policy input stack-unit all` command.

To remove only the DCB input policies applied to the specified switch, use the `no dcb-policy input stack-unit` command.

Parameters

stack-unit-id	Enter the stack unit identification.
dcb-input-policy-name	Enter the policy name for the DCB input policy.

Defaults none

Command Modes CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information The `dcb-policy input stack-unit all` command overwrites any previous `dcb-policy input stack-unit stack-unit-id` configurations. Similarly, a `dcb-policy input stack-unit stack-unit-id` command overwrites any previous `dcb-policy input stack-unit all` configuration.

Related Commands dcb-policy output stack-unit stack-ports all — applies the specified DCB output policy.

dcb-policy output

Apply the output policy with the ETS configuration to an egress interface.

Syntax

```
dcb-policy output policy-name
```

To delete the output policy, use the `no dcb-policy output` command.

Parameters	*policy-name*	Enter the output policy name.
Defaults	none	
Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

When you apply an ETS output policy to on interface, ETS-configured scheduling and bandwidth allocation take precedence over any configured settings in QoS output policies.

To remove an ETS output policy from an interface, use the `no dcb-policy output policy-name` command. ETS is enabled by default with the default ETS configuration applied (all dot1p priorities in the same group with equal bandwidth allocation).

Related Commands

dcb-output — creates a DCB output policy.

dcb-policy output stack-unit stack-ports all

Apply the specified DCB output policy on all ports of the switch stack or a single stacked switch.

Syntax

`dcb-policy output stack-unit {all | stack-unit-id} stack-ports all dcb-output-policy-name`

To remove all DCB output policies applied to the stacked ports, use the `no dcb-policy output stack-unit all` command.

To remove only the DCB output policies applied to the specified switch, use the `no dcb-policy output stack-unit` command.

Parameters	*stack-unit-id*	Enter the stack unit identification.
	dcb-output-policy-name	Enter the policy name for the DCB output policy.
Defaults	none	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

The `dcb-policy output stack-unit all` command overwrites any previous `dcb-policy output stack-unit stack-unit-id` configurations. Similarly, a `dcb-policy output stack-unit stack-unit-id` command overwrites any previous `dcb-policy output stack-unit all` configuration.

You can apply a DCB output policy with ETS configuration to all stacked ports in a switch stack or an individual stacked switch. You can apply different DCB output policies to different stack units.

Related Commands

dcb-policy input stack-unit stack-ports all — applies the specified DCB input policy.

dcb stack-unit all pfc-buffering pfc-port-count pfc-queues

Configure the PFC buffer for all switches in the stack.

Syntax

`dcb stack-unit all pfc-buffering pfc-port-count {1-56} pfc-queues {1-2}`

To remove the configuration for the PFC buffer on all switches in the stack, use the `no dcb stack-unit all pfc-buffering pfc-port-count pfc-queues` command.

Parameters

pfc-port-count {1-56}	Enter the pfc-port count. The range is 1 to 56.
pfc-queues {1-2}	Enter the pfc-queue number. The range is 1 to 2.

Defaults

The PFC buffer is enabled on all ports on the stack unit.

Command Modes

CONFIGURATION

Command History

Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

If you configure PFC on a 40GbE port, count the 40GbE port as four PFC-enabled ports in the pfc-port number you enter in the command syntax.

To achieve lossless PFC operation, the PFC port count and queue number used for the reserved buffer size that is created must be greater than or equal to the buffer size required for PFC-enabled ports and lossless queues on the switch.

You must reload the stack or a specified stack unit (use the reload command in EXEC Privilege mode) for the PFC buffer configuration to take effect.

Related Commands

dcb stack-unit pfc-buffering pfc-port pfc-queues – configures the PFC buffer for all port pipes in a specified stack unit.

dcb stack-unit pfc-buffering pfc-port-count pfc-queues

Configure the PFC buffer for all port pipes in a specified stack unit by specifying the port-pipe number, number of PFC-enabled ports, and number of configured lossless queues.

Syntax

`dcb stack-unit` *`stack-unit-id`* `[port-set` *`port-set-id`*`] pfc-buffering pfc-ports {`*`1-56`*`} pfc-queues {`*`1-2`*`}`

To remove the configuration for the PFC buffer on all port pipes in a specified stack unit, use the `no dcb stack-unit` *`stack-unit-id`* `[port-set` *`port-set-id`*`] pfc-buffering pfc-ports pfc-queues` command.

Parameters

stack-unit-id	Enter the stack unit identification. The range is from 0 to 5.
port-set	Enter the port-set identification. The only valid port-set ID (port-pipe number) on an MXL Switch is 0.
pfc-ports ***{1-56}***	Enter the pfc-ports. The range is from 1 to 56.
pfc-queues ***{1-2}***	Enter the pfc-queue number. The range is from 1 to 2.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

If you configure PFC on a 40GbE port, count the 40GbE port as four PFC-enabled ports in the pfc-port number you enter in the command syntax.

To achieve lossless PFC operation, the PFC port count and queue number used for the reserved buffer size that is created must be greater than or equal to the buffer size required for PFC-enabled ports and lossless queues on the switch.

You must reload the stack or a specified stack unit (use the `reload` command in EXEC Privilege mode) for the PFC buffer configuration to take effect.

Related Commands

dcb stack-unit pfc-buffering pfc-port pfc-queues — configures the PFC buffer for all switches in the stack.

dcbx port-role

Configure the DCBX port role the interface uses to exchange DCB information.

Syntax

`dcbx port-role {config-source | auto-downstream | auto-upstream | manual}`

To remove DCBX port role, use the `no dcbx port-role {config-source | auto-downstream | auto-upstream | manual}` command.

Parameters	**config-source | auto-downstream | auto-upstream | manual**	Enter the DCBX port role, where: • `config-source`: configures the port to serve as the configuration source on the switch. • `auto-upstream`: configures the port to receive a peer configuration. The configuration source is elected from auto-upstream ports. • `auto-downstream`: configures the port to accept the internally propagated DCB configuration from a configuration source. • `manual`: configures the port to operate only on administer-configured DCB parameters. The port does not accept a DCB configuration received form a peer or a local configuration source.
Defaults	**Manual**	
Command Modes	INTERFACE PROTOCOL LLDP	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	DCBX requires that you enable LLDP to advertise DCBX TLVs to peers. Configure DCBX operation at the INTERFACE level on a switch or globally on the switch. To verify the DCBX configuration on a port, use the `show interface dcbx detail` command.	

dcbx version

Configure the DCBX version used on the interface.

Syntax	`dcbx version {auto | cee | cin | ieee-v2.5}` To remove the DCBX version, use the `dcbx version {auto | cee | cin | ieee-v2.5}` command.	
Parameters	**auto | cee | cin | ieee-v2.5**	Enter the DCBX version type used on the interface, where: • `auto`: configures the port to operate using the DCBX version received from a peer. • `cee`: configures the port to use CDD (Intel 1.01). • `cin`: configures the port to use Cisco-Intel-Nuova (DCBX 1.0). • `ieee-v2`: configures the port to use IEEE 802.1az (Draft 2.5).

Defaults	**Auto**
Command Modes	INTERFACE PROTOCOL LLDP
Command History	**Version 9.2(0.0)** Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1 Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	DCBX requires that you enable LLDP to advertise DCBX TLVs to peers. Configure DCBX operation at the INTERFACE level on a switch or globally on the switch. To verify the DCBX configuration on a port, use the `show interface dcbx detail` command.

debug dcbx

Enable DCBX debugging.

Syntax	`debug dcbx {all | auto-detect-timer | config-exchng | fail | mgmt | resource | sem | tlv}` To disable DCBX debugging, use the `no debug dcbx` command.
Parameters	**{all | auto-detect-timer | config-exchng | fail | mgmt | resource | sem | tlv}** Enter the type of debugging, where: • `all`: enables all DCBX debugging operations. • `auto-detect-timer`: enables traces for DCBX auto-detect timers. • `config-exchng`: enables traces for DCBX configuration exchanges. • `fail`: enables traces for DCBX failures. • `mgmt`: enables traces for DCBX management frames. • `resource`: enables traces for DCBX system resource frames. • `sem`: enables traces for the DCBX state machine. • `tlv`: enables traces for DCBX TLVs.
Defaults	none
Command Modes	EXEC Privilege
Command History	**Version 9.2(0.0)** Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1 Introduced on the MXL 10/40GbE Switch IO Module.

description

Enter a text description of the DCB policy (PFC input or ETS output).

Syntax

`description text`

To remove the text description, use the `no description` command.

Parameters

text	Enter the description of the output policy. The maximum is 32 characters.

Defaults

none

Command Modes

- DCB INPUT POLICY
- DCB OUTPUT POLICY

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Related Commands

- dcb-input — creates a DCB PFC input policy.
- dcb-policy input — applies the output policy.
- dcb-output — creates a DCBETS output policy.
- dcb-policy output — applies the output policy.

ets mode on

Enable the ETS configuration so that scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBX TLV from a peer can take effect on an interface.

Syntax

`ets mode on`

To remove the ETS configuration, use the `no ets mode on` command.

Defaults

ETS mode is on.

Command Modes

DCB OUTPUT POLICY

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

If you disable ETS in an output policy applied to an interface using the `no ets mode on` command, any previously configured QoS settings at the interface or global level takes effect. If you configure QoS settings at the interface or global

level and in an output policy map (the `service-policy output` command), the QoS configuration in the output policy takes precedence.

Related Commands

- dcb-output — creates a DCB output policy.
- dcb-policy output — applies the output policy.

fcoe priority-bits

Configure the FCoE priority advertised for the FCoE protocol in application priority TLVs.

Syntax

`fcoe priority-bits` *`priority-bitmap`*

To remove the configured FCoE priority, use the `no fcoe priority-bits` command.

Parameters

priority-bitmap	Enter the priority-bitmap range. The range is from 1 to FF.

Defaults

0x8

Command Modes

PROTOCOL LLDP

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

This command is available at the global level only.

iscsi priority-bits

Configure the iSCSI priority advertised for the iSCSI protocol in application priority TLVs.

Syntax

`iscsi priority-bits` *`priority-bitmap`*

To remove the configured iSCSI priority, use the `no iscsi priority-bits` command.

Parameters

priority-bitmap	Enter the priority-bitmap range. The range is from 1 to FF.

Defaults

0x10

Command Modes

PROTOCOL LLDP

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information	This command is available at the global level only.

pfc link-delay

Configure the link delay used to pause specified priority traffic.

Syntax	`pfc link-delay value` To remove the link delay, use the `no pfc link-delay` command.	
Parameters	***value***	The range is (in quanta) from 712 to 65535. One quantum is equal to a 512-bit transmission.
Defaults	**45556** quantum	
Command Modes	DCB INPUT POLICY	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	The minimum link delay must be greater than the round-trip transmission time a peer must honor a PFC pause frame multiplied by the number of PFC-enabled ingress ports.	
Related Commands	dcb-input — creates a DCB input policy.	

pfc mode on

Enable the PFC configuration on the port so that the priorities are included in DCBX negotiation with peer PFC devices.

Syntax	`pfc mode on` To disable the PFC configuration, use the `no pfc mode on` command.	
Defaults	PFC mode is on.	
Command Modes	DCB INPUT POLICY	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	By applying a DCB input policy with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, also enable	

PFC on all DCB egress ports or configure the dot1p priority-queue assignment of PFC priorities to lossless queues (refer to `pfc no-drop queues`).

To disable PFC operation on an interface, enter the `no pfc mode on` command in DCB Input Policy Configuration mode. PFC is enabled and disabled as global DCB operation is enabled (`dcb-enable`) or disabled (`no dcb-enable`).

You cannot enable PFC and link-level flow control at the same time on an interface.

Related Commands dcb-input — creates a DCB input policy.

pfc no-drop queues

Configure the port queues that still function as no-drop queues for lossless traffic.

Syntax `pfc no-drop queues` *`queue-range`*

To remove the no-drop port queues, use the `no pfc no-drop queues` command.

Parameters

queue-range	Enter the queue range. Separate the queue values with a comma; specify a priority range with a dash; for example, `pfc no-drop queues 1,3` or `pfc no-drop queues 2-3`. The range is from 0 to 3.

Defaults No lossless queues are configured.

Command Modes INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information The maximum number of lossless queues globally supported on the switch is two.

- The following lists the dot1p priority-queue assignments.

dot1p Value in the Incoming Frame	Description heading
0	0
1	0
2	0
3	1
4	2
5	3

dot1p Value in the Incoming Frame	Description heading
6	3
7	3

pfc priority

Configure the CoS traffic to be stopped for the specified delay.

Syntax

`pfc priority` *priority-range*

To delete the pfc priority configuration, use the `no pfc priority` command.

Parameters

priority-range	Enter the 802.1p values of the frames to be paused. Separate the priority values with a comma; specify a priority range with a dash; for example, pfc priority 1,3,5-7. The range is from 0 to 7.

Defaults

none

Command Modes

DCB INPUT POLICY

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

You can enable any number of 802.1p priorities for PFC. Queues to which PFC priority traffic is mapped are lossless by default. Traffic may be interrupted due to an interface flap (going down and coming up) when you reconfigure the lossless queues for no-drop priorities in a PFC input policy and reapply the policy to an interface.

The maximum number of lossless queues supported on the I/O Aggregator switch is four.

A PFC peer must support the configured priority traffic (as DCBX detects) to apply PFC.

Related Commands

dcb-input — creates a DCB input policy.

priority-group

To use with an ETS output policy, create an ETS priority group.

Syntax

`priority-group` *group-name*

To remove the priority group, use the `no priority-group` command.

Parameters	***group-name***	Enter the name of the ETS priority group. The maximum is 32 characters.
Defaults	none	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

A priority group consists of 802.1p priority values that are grouped for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.

You must configure 802.1p priorities in priority groups associated with an ETS output policy. You can assign each dot1p priority to only one priority group.

The maximum number of priority groups supported in ETS output policies on an interface is equal to the number of data queues (4) on the port. The 802.1p priorities in a priority group can map to multiple queues.

If you configure more than one priority queue as strict priority or more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

Related Commands

- priority-list — configures the 802.1p priorities for an ETS output policy.
- set-pgid — configures the priority-group.

priority-group qos-policy

Associate the 802.1p priority traffic in a priority group with the ETS configuration in a QoS output policy.

Syntax

`priority-group` *`group-name`* `qos-policy` *`ets-policy-name`*

To remove the 802.1p priority group, use the `no priority-group qos-policy` command.

Parameters	***group-name***	Enter the group name of the 802.1p priority group. The maximum is 32 characters.
	ets-policy-name	Enter the ETS policy name.
Defaults	none	
Command Modes	DCB OUTPUT POLICY	

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

The ETS configuration associated with 802.1p priority traffic in a DCB output policy is used in DCBX negotiation with ETS peers.

If you disable ETS in an output policy applied to an interface using the `no ets mode on` command, any previously configured QoS settings at the interface or global level take effect. If you configure QoS settings at the interface or global level and in an output policy map (the `service-policy output` command), the QoS configuration in the output policy takes precedence.

Related Commands

- dcb-output — creates a DCB output policy.
- dcb-policy output — applies the output policy.

priority-list

Configure the 802.1p priorities for the traffic on which you want to apply an ETS output policy.

Syntax

`priority-list` *value*

To remove the priority list, use the `no priority-list` command.

Parameters

value	Enter the priority list value. Separate priority values with a comma; specify a priority range with a dash; for example, priority-list 3,5-7. The range is from 0 to 7.

Defaults

none

Command Modes

PRIORITY-GROUP

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

By default:

- All 802.1p priorities are grouped in priority group 0.
- 100% of the port bandwidth is assigned to priority group 0. The complete bandwidth is equally assigned to each priority class so that each class has 12 to 13%.

Related Commands

- priority-group qos-policy — associates an ETS priority group with an ETS output policy.

- set-pgid — configures the priority-group.

qos-policy-output ets

To configure the ETS bandwidth allocation and scheduling for priority traffic, create a QoS output policy.

Syntax

`qos-policy-output` *`policy-name`* `ets`

To remove the QoS output policy, use the `no qos-policy-output ets` command.

Parameters

policy-name	Enter the policy name. The maximum is 32 characters.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

If an error occurs in an ETS output-policy configuration, the configuration is ignored and the scheduler and bandwidth allocation settings are reset to the ETS default values (all priorities are in the same ETS priority group and bandwidth is allocated equally to each priority).

If an error occurs when a port receives a peer's ETS configuration, the port's configuration is reset to the previously configured ETS output policy. If no ETS output policy was previously applied, the port is reset to the default ETS parameters.

Related Commands

- scheduler — schedules the priority traffic in port queues.
- bandwidth-percentage — bandwidth percentage allocated to the priority traffic in port queues.

scheduler

Configure the method used to schedule priority traffic in port queues.

Syntax

`scheduler` *`value`*

To remove the configured priority schedule, use the `no scheduler` command.

Parameters

value	Enter schedule priority value. The valid values are: • `strict`: strict-priority traffic is serviced before any other queued traffic. • `werr`: weighted elastic round robin (werr) provides low-latency scheduling for priority traffic on port queues.

Defaults	Weighted elastic round robin (WERR) scheduling is used to queue priority traffic.	
Command Modes	POLICY-MAP-OUT-ETS	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	dot1p priority traffic on the switch is scheduled to the current queue mapping. dot1p priorities within the same queue must have the same traffic properties and scheduling method. ETS-assigned scheduling applies only to data queues, not to control queues. The configuration of bandwidth allocation and strict-queue scheduling is not supported at the same time for a priority group. If you configure both, the configured bandwidth allocation is ignored for priority-group traffic when you apply the output policy on an interface.	
Related Commands	• qos-policy-output ets — configures the ETS bandwidth allocation. • bandwidth-percentage — bandwidth percentage allocated to priority traffic in port queues.	

set-pgid

Configure the priority-group identifier.

Syntax	`set-pgid` *value* To remove the priority group, use the `no set-pgid` command.	
Parameters	***value***	Enter the priority group identification. The range is from 0 to 7.
Defaults	none	
Command Modes	PRIORITY-GROUP	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Related Commands	• priority-group qos-policy — creates an ETS priority group. • priority-list — configures the 802.1p priorities.	

show dcb

Displays the data center bridging status, the number of PFC-enabled ports, and the number of PFC-enabled queues.

Syntax

```
show dcb [stack-unit unit-number]
```

Parameters

unit number	Enter the DCB unit number. The range is from 0 to 5.

Command Modes

EXEC Privilege

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

Specify a stack-unit number on the Master switch in a stack.

Example

```
Dell# show dcb
stack-unit 0 port-set 0
      DCB Status : Enabled
    PFC Port Count : 56 (current), 56 (configured)
PFC Queue Count : 2 (current), 2 (configured)
```

show interface dcbx detail

Displays the DCBX configuration on an interface.

Syntax

```
show interface port-type slot/port dcbx detail
```

Parameters

port-type	Enter the port type.
slot/port	Enter the slot/port number.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To clear DCBX frame counters, use the `clear dcbx counters interface stack-unit/port` command.

The following describes the `show interface dcbx detail` command shown in the following example.

Field	Description
Interface	Interface type with chassis slot and port number.
Port-Role	Configured the DCBX port role: auto-upstream, auto-downstream, config-source, or manual.
DCBX Operational Status	Operational status (enabled or disabled) used to elect a configuration source and internally propagate a DCB configuration. The DCBX operational status is the combination of PFC and ETS operational status.
Configuration Source	Specifies whether the port serves as the DCBX configuration source on the switch: true (yes) or false (no).
Local DCBX Compatibility mode	DCBX version accepted in a DCB configuration as compatible. In auto-upstream mode, a port can only receive a DCBX version supported on the remote peer.
Local DCBX Configured mode	DCBX version configured on the port: CEE, CIN, IEEE v2.5, or Auto (port auto-configures to use the DCBX version received from a peer).
Peer Operating version	DCBX version that the peer uses to exchange DCB parameters.
Local DCBX TLVs Transmitted	Transmission status (enabled or disabled) of advertised DCB TLVs (see TLV code at the top of the show command output).
Local DCBX Status: DCBX Operational Version	DCBX version advertised in Control TLVs.
Local DCBX Status: DCBX Max Version Supported	Highest DCBX version supported in Control TLVs.
Local DCBX Status: Sequence Number	Sequence number transmitted in Control TLVs.
Local DCBX Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs.
Local DCBX Status: Protocol State	Current operational state of the DCBX protocol: ACK or IN-SYNC.
Peer DCBX Status: DCBX Operational Version	DCBX version advertised in Control TLVs received from the peer device.

Field	Description
Peer DCBX Status: DCBX Max Version Supported	Highest DCBX version supported in Control TLVs received from the peer device.
Peer DCBX Status: Sequence Number	Sequence number transmitted in Control TLVs received from the peer device.
Peer DCBX Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs received from the peer device.
Total DCBX Frames transmitted	Number of DCBX frames sent from the local port.
Total DCBX Frames received	Number of DCBX frames received from the remote peer port.
Total DCBX Frame errors	Number of DCBX frames with errors received.
Total DCBX Frames unrecognized	Number of unrecognizable DCBX frames received.

Example

```
Dell(conf)# show interface tengigabitethernet 0/49 dcbx detail
Dell#show interface te 0/49 dcbx detail

E-ETS Configuration TLV enabled
  e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled
  r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled
  p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled
  f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled
  i-Application Priority for iSCSI disabled
---------------------------------------------------------

Interface TenGigabitEthernet 0/49
  Remote Mac Address 00:00:00:00:00:11
   Port Role is Auto-Upstream
  DCBX Operational Status is Enabled
  Is Configuration Source? TRUE

Local DCBX Compatibility mode is CEE
  Local DCBX Configured mode is CEE
  Peer Operating version is CEE
  Local DCBX TLVs Transmitted: ErPfi

Local DCBX Status
-----------------
  DCBX Operational Version is 0
  DCBX Max Version Supported is 0
  Sequence Number: 2
  Acknowledgment Number: 2
```

```
  Protocol State: In-Sync

Peer DCBX Status:
----------------
  DCBX Operational Version is 0
  DCBX Max Version Supported is 255
  Sequence Number: 2
  Acknowledgment Number: 2
  Total DCBX Frames transmitted 27
  Total DCBX Frames received 6
  Total DCBX Frame errors 0
  Total DCBX Frames unrecognized 0
```

show interface ets

Displays the ETS configuration applied to egress traffic on an interface, including priority groups with priorities and bandwidth allocation.

Syntax

`show interface` *port-type slot/port* `ets {summary | detail}`

Parameters

***port-type slot/port* ets**	Enter the port-type slot and port ETS information.
{summary | detail}	Enter the keyword `summary` for a summary list of results or enter the keyword `detail` for a full list of results.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To clear ETS TLV counters, use the `clear ets counters interface` *port-type slot/port* command.

The following describes the `show interface summary` command shown in the following example.

Field	Description
Interface	Interface type with stack-unit and port number.
Max Supported TC Group	Maximum number of priority groups supported.
Number of Traffic Classes	Number of 802.1p priorities currently configured.
Admin mode	ETS mode: on or off. When on, the scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBX TLV from a peer can take effect on an interface.

Field	Description
Admin Parameters	ETS configuration on local port, including priority groups, assigned dot1p priorities, and bandwidth allocation.
Remote Parameters	ETS configuration on remote peer port, including admin mode (enabled if a valid TLV was received or disabled), priority groups, assigned dot1p priorities, and bandwidth allocation. If ETS admin mode is enabled on the remote port for DCBX exchange, the Willing bit received in ETS TLVs from the remote peer is included.
Local Parameters	ETS configuration on local port, including admin mode (enabled when a valid TLV is received from a peer), priority groups, assigned dot1p priorities, and bandwidth allocation.
Operational status (local port)	Port state for current operational ETS configuration: • `Init`: Local ETS configuration parameters were exchanged with the peer. • `Recommend`: Remote ETS configuration parameters were received from the peer. • `Internally propagated`: ETS configuration parameters were received from the configuration source.
ETS DCBX Oper status	Operational status of the ETS configuration on the local port: match or mismatch.
State Machine Type	Type of state machine used for DCBX exchanges of ETS parameters: Feature — for legacy DCBX versions; Asymmetric — for an IEEE version.
Conf TLV Tx Status	Status of ETS Configuration TLV advertisements: enabled or disabled.
ETS TLV Statistic: Input Conf TLV pkts	Number of ETS Configuration TLVs received.
ETS TLV Statistic: Output Conf TLV pkts	Number of ETS Configuration TLVs transmitted.
ETS TLV Statistic: Error Conf TLV pkts	Number of ETS Error Configuration TLVs received.

Example (Summary)

```
Dell(conf)# show interfaces te 0/0 ets summary
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters:
------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
```

```
0      0,1,2,3,4,5,6,7  100%      ETS
1                         0%      ETS
2                         0%      ETS
3                         0%      ETS
4                         0%      ETS
5                         0%      ETS
6                         0%      ETS
7                         0%      ETS
Priority#              Bandwidth  TSA
0                        13%      ETS
1                        13%      ETS
2                        13%      ETS
3                        13%      ETS
4                        12%      ETS
5                        12%      ETS
6                        12%      ETS
7                        12%      ETS
Remote Parameters:
-------------------
Remote is disabled
Local Parameters:
------------------
Local is enabled
TC-grp Priority#        Bandwidth TSA
0      0,1,2,3,4,5,6,7  100%      ETS
1                         0%      ETS
2                         0%      ETS
3                         0%      ETS
4                         0%      ETS
5                         0%      ETS
6                         0%      ETS
7                         0%      ETS
Priority#               Bandwidth TSA
0                        13%      ETS
1                        13%      ETS
2                        13%      ETS
3                        13%      ETS
4                        12%      ETS
5                        12%      ETS
6                        12%      ETS
7                        12%      ETS
Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
```

Example (Detail)

```
Dell(conf)# show interfaces tengigabitethernet 0/0 ets detail
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :
------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
0      0,1,2,3,4,5,6,7  100%      ETS
1                        0%        ETS
2                        0%        ETS
3                        0%        ETS
4                        0%        ETS
5                        0%        ETS
6                        0%        ETS
7                        0%        ETS
```

```
Priority#               Bandwidth TSA
0                       13%       ETS
1                       13%       ETS
2                       13%       ETS
3                       13%       ETS
4                       12%       ETS
5                       12%       ETS
6                       12%       ETS
7                       12%       ETS
Remote Parameters:
-------------------
Remote is disabled

Local Parameters :
------------------
Local is enabled
TC-grp Priority#        Bandwidth TSA
0      0,1,2,3,4,5,6,7  100%      ETS
1                       0%        ETS
2                       0%        ETS
3                       0%        ETS
4                       0%        ETS
5                       0%        ETS
6                       0%        ETS
7                       0%        ETS

Priority#               Bandwidth TSA
0                       13%       ETS
1                       13%       ETS
2                       13%       ETS
3                       13%       ETS
4                       12%       ETS
5                       12%       ETS
6                       12%       ETS
7                       12%       ETS
Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf
TLV Pkts
0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV
Pkts, 0 Error Traffic Class
TLV
Pkts
```

show interface pfc

Displays the PFC configuration applied to ingress traffic on an interface, including priorities and link delay.

Syntax

```
show interface port-type slot/port pfc {summary | detail}
```

Parameters

port-type slot/ port pfc	Enter the port-type slot and port PFC information.
{summary | detail}	Enter the keyword `summary` for a summary list of results or enter the keyword `detail` for a full list of results.

Command Modes

INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To clear the PFC TLV counters, use the `clear pfc counters interface` *`port-type slot/port`* command.

The following describes the `show interface pfc summary` command shown in the following example.

Field	Description
Interface	Interface type with stack-unit and port number.
Admin mode is on Admin is enabled	PFC admin mode is on or off with a list of the configured PFC priorities. When the PFC admin mode is on, PFC advertisements are enabled to be sent and received from peers; received PFC configuration take effect. The admin operational status for a DCBX exchange of PFC configuration is enabled or disabled.
Remote is enabled, Priority list Remote Willing Status is enabled	Operational status (enabled or disabled) of peer device for DCBX exchange of PFC configuration with a list of the configured PFC priorities. Willing status of peer device for DCBX exchange (Willing bit received in PFC TLV): enabled or disable.
Local is enabled	DCBX operational status (enabled or disabled) with a list of the configured PFC priorities.
Operational status (local port)	Port state for current operational PFC configuration: • `Init`: Local PFC configuration parameters were exchanged with the peer. • `Recommend`: Remote PFC configuration parameters were received from the peer. • `Internally propagated`: PFC configuration parameters were received from the configuration source.
PFC DCBX Oper status	Operational status for the exchange of the PFC configuration on the local port: match (up) or mismatch (down).
State Machine Type	Type of state machine used for DCBX exchanges of the PFC parameters: Feature — for legacy DCBX versions; Symmetric — for an IEEE version.
TLV Tx Status	Status of the PFC TLV advertisements: enabled or disabled.
PFC Link Delay	Link delay (in quanta) used to pause specified priority traffic.

Field	Description
Application Priority TLV: FCOE TLV Tx Status	Status of FCoE advertisements in application priority TLVs from the local DCBX port: enabled or disabled.
Application Priority TLV: SCSI TLV Tx Status	Status of ISCSI advertisements in application priority TLVs from the local DCBX port: enabled or disabled.
Application Priority TLV: Local FCOE Priority Map	Priority bitmap the local DCBX port uses in FCoE advertisements in application priority TLVs.
Application Priority TLV: Local ISCSI Priority Map	Priority bitmap the local DCBX port uses in ISCSI advertisements in application priority TLVs.
Application Priority TLV: Remote FCOE Priority Map	Status of FCoE advertisements in application priority TLVs from the remote peer port: enabled or disabled.
Application Priority TLV: Remote ISCSI Priority Map	Status of iSCSI advertisements in application priority TLVs from the remote peer port: enabled or disabled.
PFC TLV Statistics: Input TLV pkts	Number of PFC TLVs received.
PFC TLV Statistics: Output TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error pkts	Number of PFC error packets received.
PFC TLV Statistics: Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: Pause Rx pkts	Number of PFC pause frames received.

Example (Summary)

```
Dell# show interfaces tengigabitethernet 0/49 pfc summary
Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled, Priority list is 4
  Remote Willing Status is enabled
  Local is enabled
  Oper status is Recommended
  PFC DCBX Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quantams
  Application Priority TLV Parameters :
  -------------------------------------
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
```

```
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8

Dell# show interfaces tengigabitethernet 0/49 pfc detail
  Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled
  Remote Willing Status is enabled
  Local is enabled
  Oper status is recommended
  PFC DCBX Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quanta
  Application Priority TLV Parameters :
  -------------------------------------
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8
  0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts,
  0 Pause Tx pkts, 0 Pause Rx pkts
```

show interface pfc statistics

Displays counters for the PFC frames received and transmitted (by dot1p priority class) on an interface.

Syntax

```
show interface port-type slot/port pfc statistics
```

Parameters

port-type	Enter the port type.
slot/port	Enter the slot/port number.

Command Modes

INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example (Summary)

```
Dell#show interfaces te 0/3 pfc statistics
Interface TenGigabitEthernet 0/3

Priority Rx XOFF Frames Rx Total Frames Tx Total Frames
-----------------------------------------------------------
0        0              0               0
1        0              0               0
2        0              0               0
3        0              0               0
4        0              0               0
5        0              0               0
```

```
6          0                0                0
7          0                0                0
```

show qos dcb-input

Displays the PFC configuration in a DCB input policy.

Syntax

```
show qos dcb-input [pfc-profile]
```

Parameters

pfc-profile	Enter the PFC profile.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell(conf)# show qos dcb-input
dcb-input pfc-profile
  pfc link-delay 32
  pfc priority 0-1
dcb-input pfc-profile1
  no pfc mode on
  pfc priority 6-7
```

show qos dcb-output

Displays the ETS configuration in a DCB output policy.

Syntax

```
show qos dcb-output [ets-profile]
```

Parameters

[ets-profile]	Enter the ETS profile.

Command Modes

EXEC Privilege

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell# show qos dcb-output
dcb-output ets
  priority-group san qos-policy san
  priority-group ipc qos-policy ipc
  priority-group lan qos-policy lan
```

show qos priority-groups

Displays the ETS priority groups configured on the switch, including the 802.1p priority classes and ID of each group.

Syntax

```
show qos priority-groups
```

Command Modes

EXEC Privilege

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell#show qos priority-groups
priority-group ipc
  priority-list 4
  set-pgid 2
```

show stack-unit stack-ports ets details

Displays the ETS configuration applied to egress traffic on stacked ports, including ETS Operational mode on each unit and the configurated priority groups with dot1p priorities, bandwidth allocation, and scheduler type.

Syntax

```
show stack-unit {all | stack-unit} stack-ports {all | port-number} ets details
```

Parameters

stack-unit	Enter the stack unit identification.
port-number	Enter the port number.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell(conf)# show stack-unit all stack-ports all ets details

Stack unit 0 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on

Admin Parameters:
--------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
-----------------------------------------------
```

```
0      0,1,2,3,4,5,6,7  100%      ETS
1                       - -
2                       - -
3                       - -
4                       - -
5                       - -
6                       - -
7                       - -
8                       - -

Stack unit 1 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on
Admin Parameters:
--------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
-----------------------------------------------
0      0,1,2,3,4,5,6,7  100%      ETS
1                       -         -
2                       -         -
3                       -         -
4                       -         -
5                       -         -
6                       -         -
7                       -         -
8                       -         -
```

show stack-unit stack-ports pfc details

Displays the PFC configuration applied to ingress traffic on stacked ports, including PFC Operational mode on each unit with the configured priorities, link delay, and number of pause packets sent and received.

Syntax

```
show stack-unit {all | stack-unit} stack-ports {all | port-number} pfc details
```

Parameters

stack-unit	Enter the stack unit.
port-number	Enter the port number.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell(conf)# show stack-unit all stack-ports all pfc details

stack unit 0 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

```
stack unit 1 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

22

FIP Snooping

In a converged Ethernet network, an MXL Switch can operate as an intermediate Ethernet bridge to snoop on Fibre Channel over Ethernet initialization protocol (FIP) packets during the login process on Fibre Channel over Ethernet (FCoE) forwarders (FCFs).

Acting as a transit FIP snooping bridge, the switch uses dynamically-created ACLs to permit only authorized FCoE traffic to be transmitted between an FCoE end-device and an FCF. The following Dell Networking Operating System (OS) commands are used to configure and verify the FIP snooping feature.

clear fip-snooping database interface vlan

Clear FIP snooping information on a VLAN for a specified FCoE MAC address, ENode MAC address, or FCF MAC address, and remove the corresponding ACLs FIP snooping generates.

Syntax

```
clear fip-snooping database interface vlan vlan-id {fcoe-mac-address | enode-mac-address | fcf-mac-address}
```

Parameters

fcoe-mac-address	Enter the FCoE MAC address to be cleared of FIP snooping information.
enode-mac-address	Enter the ENode MAC address to be cleared of FIP snooping information.
fcf-mac-address	Enter the FCF MAC address to be cleared of FIP snooping information.

Command Modes EXEC Privilege

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

clear fip-snooping statistics

Clears the statistics on the FIP packets snooped on all VLANs, a specified VLAN, or a specified port interface.

Syntax

```
clear fip-snooping statistics [interface vlan vlan-id | interface port-type port/slot | interface port-channel port-channel-number]
```

Parameters

vlan-id	Enter the VLAN ID of the FIP packet statistics to be cleared.

	port-type port/ slot	Enter the port-type and slot number of the FIP packet statistics to be cleared.
	port-channel-number	Enter the port channel number of the FIP packet statistics to be cleared.
Command Modes	EXEC Privilege	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

feature fip-snooping

Enable FCoE transit and FIP snooping on a switch.

Syntax	`feature fip-snooping` To disable the FCoE transit feature, use the `no feature fip-snooping` command.
Defaults	Disabled
Command Modes	CONFIGURATION

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

fip-snooping enable

Enable FIP snooping on all VLANs or on a specified VLAN.

Syntax	`fip-snooping enable` To disable the FIP snooping feature on all or a specified VLAN, use the `no fip-snooping enable` command.
Defaults	FIP snooping is disabled on all VLANs.
Command Modes	• CONFIGURATION • VLAN INTERFACE

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

The maximum number of FCFs supported per FIP snooping-enabled VLAN is four.
The maximum number of FIP snooping sessions supported per ENode server is 16.

fip-snooping fc-map

Configure the FC-MAP value FIP snooping uses on all VLANs.

Syntax

`fip-snooping fc-map` *`fc-map-value`*

To return the configured FM-MAP value to the default value, use the `no fip-snooping fc-map` command.

Parameters

fc-map-value	Enter the FC-MAP value FIP snooping uses. The range is from 0EFC00 to 0EFCFF.

Defaults

0x0EFC00

Command Modes

- CONFIGURATION
- VLAN INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

fip-snooping port-mode fcf

Configure the port for bridge-to-FCF links.

Syntax

`fip-snooping port-mode fcf`

To disable the bridge-to-FCF link on a port, use the `no fip-snooping port-mode fcf` command.

Command Modes

INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

The maximum number of FCFs supported per FIP snooping-enabled VLAN is four.

23

High Availability (HA)

High availability (HA) in the Dell Networking operating software (FTOS) is configuration synchronization to minimize recovery time in the event of a route processor module (RPM) failure. The feature is available on the S4810 S4820T platform.

In general, a protocol is defined as "hitless" in the context of an RPM failure/failover and not failures of a line card, SFM, or power module. A protocol is defined as hitless if an RPM failover has no impact on the protocol.

You must specifically enable some protocols for HA. Some protocols are only hitless if related protocols are also enabled as hitless (for example, the `redundancy protocol` command).

redundancy force-failover

Force the secondary stack unit to become the primary stack unit. You can also use this command to upgrade the software on one stack unit from the other when the other has been loaded with the upgraded software.

Z9000 S4810 S4820T

Syntax	`redundancy force-failover {stack-unit` *`unit-number`*`}`	
Parameters	**stack-unit *unit-number***	Enter the keyword `stack-unit` then the stack-unit ID number. The range is from 0 to 7.
Default	Not configured.	
Command Modes	EXEC Privilege	
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.	
	The following is a list of the FTOS version history for this command.	
	Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.19.0	Introduced on the S4820T.
	Version 8.3.12.0	Introduced on the S4810.
	Version 8.1.1.0	Introduced on the E-Series ExaScale.
	Version 7.5.1.0	Introduced on the C-Series.

Version 7.6.1.0	Introduced on the E-Series.

Usage Information

To provide a hitless or warm upgrade, use this command. A hitless upgrade means that a software upgrade does not require a reboot of the line cards. A warm upgrade means that a software upgrade requires a reset of the line cards. A warm upgrade is possible for major releases and lower, while a hitless upgrade can only support patch releases.

show redundancy

Display the current redundancy configuration.

Z9000 S4810 S4820T

Syntax

```
show redundancy
```

Command Modes

- EXEC
- EXEC Privilege

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 8.1.1.0	Introduced on the E-Series ExaScale.
Version 7.5.1.0	Introduced on the C-Series.
Version 7.6.1.0	Introduced on the E-Series.

Usage Information

The following describes the `show redundancy` command shown in the following example.

Field	Description
RPM Status	Displays the following information: • Slot number of the RPM. • Whether the RPM is Primary or Standby. • The state of the RPM: Active, Standby, Booting, or Offline.

Field	Description
	• Whether the link to the second RPM is up or down.
PEER RPM Status	Displays the state of the second RPM, if present
RPM Redundancy Configuration	Displays the following information: • which RPM is the preferred Primary on next boot (the `redundancy primary` command) • the data sync method configured (the `redundancy synchronize` command) • the failover type (you cannot change this type; it is software-dependent). **Hot Failover** means that the running configuration and routing table are applied on secondary RPM. **Fast Failover** means that the running configuration is not applied on the secondary RPM until failover occurs, and the routing table on line cards is cleared during failover. • the status of auto booting the RPM (the `redundancy disable-auto-reboot` command) • the parameter for auto failover limit control (the `redundancy auto-failover-limit` command) •
RPM Failover Record	Displays the following information: • RPM failover counter (to reset the counter, use the `redundancy reset-counter` command) • the time and date of the last RPM failover • the reason for the last RPM failover
Last Data Sync Record	Displays the data sync information and the timestamp for the data sync: • Start-up Config is the contents of the startup-config file. • Line Card Config is the line card types configured and interfaces on those line cards. • Runtime Event Log is the contents of the Event log. • Running Config is the current running-config. This field only appears when you enter the command from the Primary RPM.

Example (S4820T)

```
FTOS#show redundancy
-- Stack-unit Status --
-----------------------------------------------
Mgmt ID:                    0
Stack-unit ID:              0
Stack-unit Redundancy Role: Primary
Stack-unit State:           Active
Stack-unit SW Version:      FIT-R2D2-1-0-0-89
Link to Peer:               Down
Peer Stack-unit:            not present
```

```
-- Stack-unit Redundancy Configuration --
-----------------------------------------------
Primary Stack-unit:         mgmt-id 0
Auto Data Sync:             Full
Failover Type:              Hot Failover
Auto reboot Stack-unit: Disabled
Auto failover limit:    3 times in 60 minutes

-- Stack-unit Failover Record --
-----------------------------------------------
Failover Count:          0
Last failover timestamp: None
Last failover Reason:    None
Last failover type:      None

-- Last Data Block Sync Record: --
-----------------------------------------------
Stack Unit Config:       no block sync done
Start-up Config:         no block sync done
Runtime Event Log:       no block sync done
Running Config:          no block sync done
ACL Mgr:                 no block sync done
LACP:                    no block sync done
STP:                     no block sync done
SPAN:                    no block sync done

FTOS#
```

Example

```
FTOS#show redundancy

-- RPM Status --
-----------------------------------------------
RPM Slot ID:         1
RPM Redundancy Role: Primary
RPM State:           Active
RPM SW Version:      7.5.1.0
Link to Peer:        Up

-- PEER RPM Status --
-----------------------------------------------
RPM State:       Standby
RPM SW Version:  7.5.1.0

-- RPM Redundancy Configuration --
-----------------------------------------------
Primary RPM:         rpm0
Auto Data Sync:      Full
Failover Type:       Hot Failover
Auto reboot RPM:     Enabled
Auto failover limit: 3 times in 60 minutes

-- RPM Failover Record --
-----------------------------------------------
Failover Count:           1
Last failover timestamp:  Jul 13 2007 21:25:32
Last failover Reason:     User request

-- Last Data Block Sync Record: --
-----------------------------------------------
Line Card Config:  succeeded Jul 13 2007 21:28:53
Start-up Config:   succeeded Jul 13 2007 21:28:53
SFM Config State:  succeeded Jul 13 2007 21:28:53
```

```
Runtime Event Log: succeeded Jul 13 2007 21:28:53
Running Config:    succeeded Jul 13 2007 21:28:53
FTOS#
```

24

iSCSI Optimization

Internet small computer system interface (iSCSI) optimization enables quality-of-service (QoS) treatment for iSCSI storage traffic.

To configure and verify the iSCSI optimization feature, use the following Dell Networking operating software commands.

advertise dcbx-app-tlv

Configure DCBX to send iSCSI TLV advertisements.

Syntax	`advertise dcbx-app-tlv iscsi` To disable DCBX iSCSI TLV advertisements, use the `no advertise dcbx-app-tlv iscsi` command.	
Defaults	Disabled.	
Command Modes	PROTOCOL LLDP	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	You can configure iSCSI TLVs to send either globally or on a specified interface. The interface configuration takes priority over global configuration.	

iscsi aging time

Set the aging time for iSCSI sessions.

Syntax	`iscsi aging time` *time* To remove the iSCSI session aging time, use the `no iscsi aging time` command.	
Parameters	***time***	Enter the aging time for the iSCSI session. The range is from 5 to 43,200 minutes.
Defaults	**10 minutes**	
Command Modes	CONFIGURATION	

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

iscsi cos

Set the QoS policy that is applied to the iSCSI flows.

Syntax

```
iscsi cos {enable | disable | dot1p vlan-priority-value
[remark] | dscp dscp-value [remark]}
```

To disable the QoS policy, use the `no iscsi cos dscp` command.

Parameters

enable	Enter the keyword `enable` to allow the application of preferential QoS treatment to iSCSI traffic so that the iSCSI packets are scheduled in the switch with a dot1p priority 4 regardless of the VLAN priority tag in the packet. The default is: the iSCSI packets are handled with dotp1 priority 4 without remark.
disable	Enter the keyword `disable` to disable the application of preferential QoS treatment to iSCSI frames.
dot1p *vlan-priority-value*	Enter the dot1p value of the VLAN priority tag assigned to the incoming packets in an iSCSI session. The range is from 0 to 7. The default is the dot1p value in ingress iSCSI frames is not changed and is the same priority is used in iSCSI TLV advertisements if you did not enter the `iscsi priority-bits` command.
dscp *dscp-value*	Enter the DSCP value assigned to the incoming packets in an iSCSI session. The valid range is from 0 to 63. The default is: the DSCP value in ingress packets is not changed.
remark	Marks the incoming iSCSI packets with the configured dot1p or DSCP value when they egress to the switch. The default is: the dot1and DSCP values in egress packets are not changed.

Defaults The default dot1p VLAN priority value is 4 without the `remark` option.

Command Modes CONFIGURATION

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

iscsi enable

Globally enable iSCSI optimization.

Syntax	`iscsi enable` To disable iSCSI optimization, use the `no iscsi enable` command.	
Parameters	***enable***	Enter the keyword `enable` to enable the iSCSI optimization feature.
Defaults	Disabled.	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	When you enable the iSCSI feature using the `iscsi enable` command, flow control settings are set to `rx on tx off` on all interfaces.	

iscsi priority-bits

Configure the priority bitmap that advertises in the iSCSI application TLVs.

Syntax	`iscsi priority-bits` To remove the configured priority bitmap, use the `no iscsi priority-bits` command.	
Defaults	**4** (0x10 in the bitmap)	
Command Modes	PROTOCOL LLDP (only on the global, not on the interface)	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

iscsi profile-compellant

Configure the auto-detection of Dell Compellent arrays on a port.

Syntax	`iscsi profile-compellent`
Defaults	Dell Compellent disk arrays are not detected.

Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

iscsi target port

Configure the iSCSI target ports and optionally, the IP addresses on which iSCSI communication is monitored.

Syntax

`iscsi target port [`*`tcp-port-2...tcp-port-16`*`]ip-address [`*`ip-address`*`]`

To remove the configured iSCSI target ports or IP addresses, use the `no iscsi target port` command.

Parameters

tcp-port-2...tcpport - 16	Enter the tcp-port number of the iSCSI target ports. The `tcp-port-n` is the TCP port number or a list of TCP port numbers on which the iSCSI target listens to requests. Separate port numbers with a comma. The default is **860, 3260**.
ip-address	(Optional) Enter the ip-address that the iSCSI monitors. The ip-address specifies the IP address of the iSCSI target.

Defaults	**860, 3260**	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

You can configure up to 16 target TCP ports on the switch in one command or multiple commands.

When you use the `no iscsi target port` command and the TCP port you wish to delete is one bound to a specific IP address, the IP address value must be included in the command.

iSCSI Optimization Prerequisites

The following are iSCSI optimization prerequisites.

- iSCSI optimization requires LLDP on the switch. LLDP is enabled by default (refer to Link Layer Discovery Protocol (LLDP)).

- iSCSI optimization requires configuring two ingress ACL groups The ACL groups are allocated after iSCSI Optimization is configured. (refer to When to Use CAM Profiling).

Configuring iSCSI Optimization

To configure iSCSI optimization, use the following commands.

1. **For a non-DCB environment:** Enable session monitoring.

   CONFIGURATION mode

   ```
   cam-acl l2acl 4 ipv4acl 4 ipv6acl 0 ipv4qos 2 l2qos 1 l2pt 0 ipmacacl 0
   vman-qos 0 ecfmacl 0 fcoeacl 0 iscsioptacl 2
   ```

   NOTE: In FTOS Version 9.2.(0.0), content addressable memory (CAM) allocation is optional. If CAM is not allocated, the following features are disabled:

   - session monitoring
   - aging
   - class of service

   You can enable iSCSI even when allocated with zero (0) CAM blocks. However, if no CAM blocks are allocated, session monitoring is disabled and this information displays in the `show iscsi` command.

2. **For a non-DCB environment:** Enable iSCSI.

   CONFIGURATION mode

   ```
   iscsi enable
   ```

3. **For a DCB environment:** Configure iSCSI Optimization.

   EXEC Privilege mode

   iSCSI configuration: `copy CONFIG_TEMPLATE/iSCSI_DCB_Config running-config`.

   The configuration files are stored in the flash memory in the CONFIG_TEMPLATE file.

   NOTE: DCB/DCBx is enabled when you apply the iSCSI configuration in step 3. If you manually apply the iSCSI configuration by following steps 1 and 2, enable link layer discovery protocol (LLDP) before enabling iSCSI in step 2. You cannot disable LLDP if you enable iSCSI.

4. Save the configuration on the switch.

   EXEC Privilege mode

   ```
   write memory
   ```

5. Reload the switch.

   EXEC Privilege mode

   ```
   reload
   ```

   After the switch is reloaded, DCB/ DCBx and iSCSI monitoring are enabled.

6. (Optional) Configure the iSCSI target ports and optionally the IP addresses on which iSCSI communication is monitored.

   CONFIGURATION mode

   `[no] iscsi target port tcp-port-1 [tcp-port-2...tcp-port-16] [ip-address address]`

   - *tcp-port-n* is the TCP port number or a list of TCP port numbers on which the iSCSI target listens to requests. You can configure up to 16 target TCP ports on the switch in one command or multiple commands. The default is **860**, **3260**.

     Separate port numbers with a comma. If multiple IP addresses are mapped to a single TCP port, use the `no iscsi target port tcp-port-n` command to remove all IP addresses assigned to the TCP number.

     To delete a specific IP address from the TCP port, use the `no iscsi target port tcp-port-n ip-address address` command to specify the address to be deleted.
   - *ip-address* specifies the IP address of the iSCSI target. When you enter the `no` form of the command, and the TCP port you want to delete is one bound to a specific IP address, include the IP address value in the command.

     If multiple IP addresses are mapped to a single TCP port, use the `no iscsi target port` command to remove all IP addresses assigned to the TCP port number.

     To remove a single IP address from the TCP port, use the `no iscsi target port ip-address` command.

7. (Optional) Set the QoS policy that is applied to the iSCSI flows.

   CONFIGURATION mode

   `[no] iscsi cos {enable | disable | dot1p vlan-priority-value [remark] | dscp dscp-value [remark]}`

   - `enable`: enables the application of preferential QoS treatment to iSCSI traffic so that iSCSI packets are scheduled in the switch with a dot1p priority 4 regardless of the VLAN priority tag in the packet. The default is: iSCSI packets are handled with dotp1 priority 4 without `remark`.
   - `disable`: disables the application of preferential QoS treatment to iSCSI frames.
   - `dot1p vlan-priority-value`: specifies the virtual local area network (VLAN) priority tag assigned to incoming packets in an iSCSI session. The range is from 0 to 7. The default is: the dot1p value in ingress iSCSI frames is not changed and the same priority is used in iSCSI TLV advertisements if you do not enter the `iscsi priority-bits` command (Step 10).
   - `dscp dscp-value`: specifies the DSCP value assigned to incoming packets in an iSCSI session. The range is from 0 to 63. The default is: the DSCP value in ingress packets is not changed.
   - `remark`: marks incoming iSCSI packets with the configured dot1p or DSCP value when they egress the switch. The default is: the dot1 and DSCP values in egress packets are not changed.

8. (Optional) Set the aging time for iSCSI session monitoring.

   CONFIGURATION mode

   `[no] iscsi aging time time`.

   The range is from 5 to 43,200 minutes.

   The default is **10 minutes**.

9. (Optional) Configures DCBX to send iSCSI TLV advertisements.

   LLDP CONFIGURATION mode or INTERFACE LLDP CONFIGURATION mode

   `[no] advertise dcbx-app-tlv iscsi`.

   You can send iSCSI TLVs either globally or on a specified interface. The interface configuration takes priority over global configuration.

   The default is **Enabled**.

10. (Optional) Configures the advertised priority bitmap in iSCSI application TLVs.
    LLDP CONFIGURATION mode
    `[no] iscsi priority-bits`.
    The default is **4** (0x10 in the bitmap).
11. (Optional) Configures the auto-detection of Compellent arrays on a port.
    INTERFACE mode
    `[no] iscsi profile-compellent`.
    The default is: Compellent disk arrays are not detected.

25

Interfaces

The commands in this chapter are supported by Dell Networking operating software (Dell).
This chapter contains the following sections:

- Basic Interface Commands
- Port Channel Commands
- Time Domain Reflectometer (TDR)
- UDP Broadcast

Basic Interface Commands

The following commands are for Physical, Loopback, and Null interfaces.

clear counters

Clear the counters used in the `show interfaces` commands for all virtual router redundancy protocol (VRRP) groups, virtual local area networks (VLANs), and physical interfaces, or selected ones.

Syntax

```
clear counters [interface] [vrrp [{vrid | vrf instance}]|
learning-limit]
```

Parameters

interface (OPTIONAL) Enter any of the following keywords and slot/port or number to clear counters from a specified interface:

- For a Loopback interface, enter the keyword `loopback` then a number from 0 to 16383.
- For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128.
- For the management interface on the RPM, enter the keyword `ManagementEthernet` then slot/port information. The slot range is from 0 to 1 and the port range is 0.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For a VLAN, enter the keyword `VLAN` then a number from 1 to 4094.

vrrp *vrid*] (OPTIONAL) Enter the keyword `vrrp` to clear the counters of all VRRP groups. To clear the counters of a specified group, enter a VRID number from 1 to 255.

	vrrp [vrf *instance*]	(OPTIONAL) Enter the keyword `vrrp` to clear the counters of all VRRP groups. To clear the counters of VRRP groups in a specified VRF instance, enter the name of the instance (32 characters maximum).
	learning-limit	(OPTIONAL) Enter the keywords `learning-limit` to clear unknown source address (SA) drop counters when MAC learning limit is configured on the interface.

Defaults Without an interface specified, the command clears all interface counters.

Command Modes EXEC Privilege

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell#clear counters
Clear counters on all interfaces [confirm]
```

Related Commands mac learning-limit — allows aging of MACs even though a learning-limit is configured or disallow station move on learned MACs.

show interfaces — displays information on the interfaces.

description

Assign a descriptive text string to the interface.

Syntax `description` *desc_text*

To delete a description, use the `no description` command.

Parameters

desc_text	Enter a text string up to 240 characters long. To use special characters as a part of the description string, you must enclose the whole string in double quotes.

Defaults none

Command Modes INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information	**Important Points to Remember:** • Spaces between characters are not preserved after entering this command unless you enclose the entire description in quotation marks (“*desc_text*”). • Entering a text string after the `description` command overwrites any previous text string that you previously configured as the description. • The `shutdown` and `description` commands are the only commands that you can configure on an interface that is a member of a port-channel. • Use the `show interfaces description` command to display descriptions configured for each interface.
Related Commands	show interfaces description — displays the description field of the interfaces.

flowcontrol

Control how the system responds to and generates 802.3x pause frames on 10G and 40Gig stack units.

Syntax `flowcontrol rx {off | on} tx {off | on} threshold`

Parameters

rx on	Enter the keywords `rx on` to process the received flow control frames on this port. This is the default value for the receive side.
rx off	Enter the keywords `rx off` to ignore the received flow control frames on this port.
tx on	Enter the keywords `tx on` to send control frames from this port to the connected device when a higher rate of traffic is received. This is the default value on the send side.
tx off	Enter the keywords `tx off` so that flow control frames are not sent from this port to the connected device when a higher rate of traffic is received.

Defaults

- `rx off`
- `tx off`

Command Modes INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information The globally assigned 48-bit Multicast address 01-80-C2-00-00-01 is used to send and receive pause frames. To allow full-duplex flow control, stations implementing the pause operation instruct the MAC to enable the reception of frames with a destination address equal to this multicast address.

The pause:

- Starts when *either* the packet pointer or the buffer threshold is met (whichever is met first). When the discard threshold is met, packets are dropped.
- Ends when *both* the packet pointer and the buffer threshold fall below 50% of the threshold settings.

The *discard threshold* defines when the interface starts dropping the packet on the interface. This may be necessary when a connected device does not honor the flow control frame sent by the switch. The discard threshold should be larger than the *buffer threshold* so that the buffer holds at least hold at least three packets.

On 4–port 10G stack units: Changes in the flow-control values may not be reflected automatically in the `show interface` output for 10G interfaces. This is because 10G interfaces do not support auto-negotiation.

Important Points to Remember

- Do not enable `tx pause` when buffer carving is enabled. For information and assistance, consult Dell Networking TAC.
- Asymmetric flow control (`rx on tx off`, or `rx off tx on`) setting for the interface port less than 100 Mb/s speed is not permitted. The following error is returned:

  ```
  Can't configure Asymmetric flowcontrol when speed <1G,
  config ignored
  ```

- The only configuration applicable to half duplex ports is `rx off tx off`. The following error is returned:

  ```
  Cannot configure Asymmetric flowcontrol when speed <1G,
  config ignored>
  ```

- You cannot configure half duplex when the flow control configuration is on (default is `rx on tx on`). The following error is returned: `Cannot configure half duplex when flowcontrol is on, config ignored`

NOTE: The flow control must be off (`rx off tx off`) before configuring the half duplex.

Example (partial)

```
Dell(conf-if-tengig-0/1)#show config
!
interface TenGigabitEthernet 0/1
no ip address
switchport
no negotiation auto
flowcontrol rx off tx on
no shutdown
...
```

Example (Values)

This Example shows how the Dell Networking OS negotiates the flow control values between two Dell Networking chassis connected back-to-back using 1G copper ports.

```
Configured

LocRxConf LocTxConf RemoteRxConf RemoteTxConf
off       off        off          off
                     off          on
                     on           off
                     on           on
```

```
off        on          off           off
                       off           on
                       on            off
                       on            on

on         off         off           off
                       off           on
                       on            off
                       on            on

on         on          off           off
                       off           on
                       on            off
                       on            on

LocNegRx LocNegTx RemNegRx RemNegTx
off      off      off      off
off      off      off      off
off      off      off      off
off      off      off      off

off      off      off      off
off      off      off      off
off      on       on       off
off      off      off      off

off      off      off      off
on       off      off      on
on       on       on       on
on       on       on       on

off      off       off     off
off      off       off     off
on       on        on      on
on       on        on      on
```

Related Commands

show running-config — displays the flow configuration parameters (non-default values only).

show interfaces — displays the negotiated flow control parameters.

interface

Configure a physical interface on the switch.

Syntax

```
interface interface
```

Parameters

interface Enter one of the following keywords and slot/port or number information:

- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For a Fibre Channel interface, enter the keyword `FibreChannel`, then the slot/port information.

Defaults	Not configured.	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	You cannot delete a physical interface. By default, physical interfaces are disabled (`shutdown`) and are in Layer 3 mode. To place an interface in mode, ensure that the interface's configuration does not contain an IP address and enter the `Port Channel Commands` command.	
Example	`Dell(conf)#interface tengig 0/0` `Dell(conf-if-tengig-0/0)#exit#`	
Related Commands	interface loopback — configures a Loopback interface. interface null — configures a Null interface. interface port-channel — configures a port channel. interface vlan — configures a VLAN. show interfaces — displays the interface configuration.	

interface ManagementEthernet

Configure the Management port on the system.

Syntax	`interface ManagementEthernet` *slot/port*	
Parameters	***slot/port***	Enter the keyword `ManagementEthernet`, then the slot number (0 or 1) and port number zero (0).
Defaults	Not configured.	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	You cannot delete a Management port. The Management port is enabled by default (`no shutdown`). To assign an IP address to the Management port, use the `ip address` command.	

Example

```
Dell(conf)#interface managementethernet 0/0
Dell(conf-if-ma-0/0)#
```

Related Commands

management route — configures a static route that points to the Management interface or a forwarding router.

duplex (1000/10000 Interfaces)— configure duplex mode on any physical interfaces where the speed is set to 1000/10000.

interface range

This command permits configuration of a range of interfaces to which subsequent commands are applied (bulk configuration). Using the `interface range` command, you can enter identical commands for a range of interface.

Syntax

`interface range` *`interface`*, *`interface`*,...

Parameters

interface, interface,...

Enter the keywords `interface range` and one of the interfaces — slot/port, port-channel, or VLAN number. Select the range of interfaces for bulk configuration. You can enter up to six comma-separated ranges. Spaces are not required between the commas. Comma-separated ranges can include VLANs, port-channels, and physical interfaces.

Slot/Port information must contain a space before and after the dash. For example, `interface range gigabitethernet 0/1 - 5` is valid; `interface range gigabitethernet 0/1-5` is NOT valid.

- For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For a VLAN interface, enter the keyword `vlan` then a number from 1 to 4094.

Defaults

none

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

When creating an interface range, interfaces appear in the order they are entered; they are not sorted. The command verifies that interfaces are present (physical) or configured (logical).

Important Points to Remember:

- Bulk configuration is created if at least one interface is valid.
- Non-existing interfaces are excluded from the bulk configuration with a warning message.
- The `interface range` prompt includes interface types with slot/port information for valid interfaces. The prompt allows for a maximum of 32 characters. If the bulk configuration exceeds 32 characters, it is represented by an ellipsis (...).
- When the `interface range` prompt has multiple port ranges, the smaller port range is excluded from the prompt.
- If overlapping port ranges are specified, the port range is extended to the smallest start port and the biggest end port.

Example (Bulk)

```
Dell(conf)#interface range so 2/0-1, te 10/0, gi 3/0, fa 0/0
% Warning: Non-existing ports (not configured) are ignored by
interface-range
```

Example (Multiple Ports)

```
Dell(conf)#interface range gi 2/0 - 23, gi 2/1 - 10
Dell(conf-if-range-gi-2/0-23#
```

Example (Overlapping Ports)

```
Dell(conf)#interface range gi 2/1 - 11, gi 2/1 - 23
Dell(conf-if-range-gi-2/1-23#
```

Usage Information

Only VLAN and port-channel interfaces created using the `interface vlan` and `interface port-channel` commands can be used in the `interface range` command.

Use the `show running-config` command to display the VLAN and port-channel interfaces. VLAN or port-channel interfaces that are not displayed in the `show running-config` command cannot be used with the bulk configuration feature of the `interface range` command. You cannot create virtual interfaces (VLAN, Port-channel) using the `interface range` command.

NOTE: If a range has VLAN, physical, port-channel, and SONET interfaces, only commands related to physical interfaces can be bulk configured. To configure commands specific to VLAN or port-channel, only those respective interfaces should be configured in a particular range.

Example (Single Range)

This example shows a single range bulk configuration.

```
Dell(config)# interface range gigabitethernet 5/1 - 23
Dell(config-if-range)# no shutdown
Dell(config-if-range)#
```

Example (Multiple Range)

This example shows how to use commas to add different interface types to the range enabling all Gigabit Ethernet interfaces in the range 5/1 to 5/23 and both Ten-Gigabit Ethernet interfaces 1/1 and 1/2.

```
Dell(config-if)# interface range gigabitethernet5/1-23,
tengigabitethernet1/1-2
Dell(config-if-range)# no shutdown
Dell(config-if-range)#
```

Example (Multiple Range)

This example shows how to use commas to add SONET, VLAN, and port-channel interfaces to the range.

```
Dell(config-if)# interface range gigabitethernet5/1-23,
tengigabitethernet1/1-2,
Vlan 2-100, Port 1-25
Dell(config-if-range)# no shutdown
Dell(config-if-range)#
```

Related Commands

interface port-channel — configures a port channel group.

interface vlan — configures a VLAN interface.

show config (from INTERFACE RANGE mode) — shows the bulk configuration interfaces.

show range — shows the bulk configuration ranges.

interface range macro (define) — defines a macro for an interface-range.

interface vlan

Configure a VLAN. You can configure up to 4096 VLANs.

Syntax

`interface vlan` *`vlan-id`*

To delete a VLAN, use the `no interface vlan` *`vlan-id`* command.

Parameters

vlan-id	Enter a number as the VLAN Identifier. The range is from 1 to 4096.

Defaults

Not configured, except for the Default VLAN, which is configured as VLAN 1.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

For more information about VLANs and the commands to configure them, refer to the Virtual LAN (VLAN) Commands.

FTP, TFTP, and SNMP operations are not supported on a VLAN. MAC ACLs are not supported in VLANs. IP ACLs are supported. For more information, refer to the Access Control Lists (ACL) chapter.

Example

```
Dell(conf)#int vlan 3
Dell(conf-if-vl-3)#
```

Related Commands

interface — configures a physical interface.

interface loopback — configures a loopback interface.

interface null — configures a null interface.

interface port-channel — configures a port channel group.

show vlan — displays the current VLAN configuration on the switch.

shutdown — disables/enables the VLAN.

tagged — adds a Layer 2 interface to a VLAN as a tagged interface.

untagged — adds a Layer 2 interface to a VLAN as an untagged interface.

keepalive

Send keepalive packets periodically to keep an interface alive when it is not transmitting data.

Syntax	`keepalive [seconds]` To stop sending keepalive packets, use the `no keepalive` command.	
Parameters	***seconds***	(OPTIONAL) For interfaces with PPP encapsulation enabled, enter the number of seconds between keepalive packets. The range is from 0 to 23767. The default is **10 seconds**.
Defaults	Enabled.	
Command Modes	INTERFACE	
Command History	**Version 8.3.16.1**	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	When you configure `keepalive`, the system sends a self-addressed packet out of the configured interface to verify that the far end of a WAN link is up. When you configure `no keepalive`, the system does not send keepalive packets and so the local end of a WAN link remains up even if the remote end is down.	

mtu

Set the link maximum transmission unit (MTU) (frame size) for an Ethernet interface.

Syntax	`mtu value` To return to the default MTU value, use the `no mtu` command.	
Parameters	***value***	Enter a maximum frame size in bytes. The range is from 594 to 9252. MXL Switch Range is from 594 to 12000. The default is **1554**.
Defaults	**1554**	
Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

If the packet includes a Layer 2 header, the difference between the link MTU and IP MTU (`ip mtu` command) must be enough bytes to include the Layer 2 header.

- The IP MTU is adjusted automatically when you configure the Layer 2 MTU with the `mtu` command.

When you enter the `no mtu` command, The Dell Networking OS reduces the IP MTU value to 1536 bytes.

Link MTU and IP MTU considerations for port channels and VLANs are as follows.

port channels:

- All members must have the same link MTU value and the same IP MTU value.
- The port channel link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the channel members. For example, if the members have a link MTU of 2100 and an IP MTU 2000, the port channel's MTU values cannot be higher than 2100 for link MTU or 2000 bytes for IP MTU.

VLANs:

- All members of a VLAN must have same IP MTU value.
- Members can have different Link MTU values. Tagged members must have a link MTU 4 bytes higher than untagged members to account for the packet tag.
- The VLAN link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the VLAN members. For example, the VLAN contains tagged members with Link MTU of 1522 and IP MTU of 1500 and untagged members with Link MTU of 1518 and IP MTU of 1500. The VLAN's Link MTU cannot be higher than 1518 bytes and its IP MTU cannot be higher than 1500 bytes.

The following shows the difference between Link MTU and IP MTU.

Layer 2 Overhead	Link MTU and IP MTU Delta
Ethernet (untagged)	18 bytes
VLAN Tag	22 bytes
Untagged Packet with VLAN-Stack Header	22 bytes
Tagged Packet with VLAN-Stack Header	26 bytes

negotiation auto

Enable auto-negotiation on an interface.

Syntax

`negotiation auto`

To disable auto-negotiation, use the `no negotiation auto` command.

Defaults

Enabled.

Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

The `no negotiation auto` command is only available if you first manually set the speed of a port to **10Mbits** or **100Mbits**.

The `negotiation auto` command provides a `mode` option for configuring an individual port to forced-master/forced slave after you enable auto-negotiation.

If you do not use the `mode` option, the default setting is **slave**. If you do not configure forced-master or forced-slave on a port, the port negotiates to either a master or a slave state. Port status is one of the following:

- Forced-master
- Force-slave
- Master
- Slave
- Auto-neg Error — typically indicates that both ends of the node are configured with forced-master or forced-slave.

△ **CAUTION: Ensure that one end of your node is configured as forced-master and one is configured as forced-slave. If both are configured the same (that is, forced-master or forced-slave), the `show interfaces` command flaps between an auto-neg-error and forced-master/slave states.**

You can display master/slave settings with the `show interfaces` command.

Example (Master/Slave)

```
Dell(conf)# int tengig 0/0
Dell(conf-if)#neg auto
Dell(conf-if-autoneg)# ?

end         Exit from configuration mode
exit        Exit from autoneg configuration mode
mode        Specify autoneg mode
no          Negate a command or set its defaults
show        Show autoneg configuration information
Dell(conf-if-autoneg)#mode ?
forced-master Force port to master mode
forced-slave Force port to slave mode
Dell(conf-if-autoneg)#
```

Example (Master/Slave, partial)

```
Dell#show interfaces configured
TenGigabitEthernet 13/18 is up, line protocol is up
Hardware is Dell Force10Eth, address is 00:01:e8:05:f7:fc
   Current address is 00:01:e8:05:f7:fc
Interface index is 474791997
Internet address is 1.1.1.1/24
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit, Mode full duplex, Master
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interfaces" counters 00:12:42
Queueing strategy: fifo
```

```
Input Statistics:
...
```

User Information

Both sides of the link must have auto-negotiation enabled or disabled for the link to come up.

The following details the possible speed and auto-negotiation combinations for a line between two 10/100/1000 Base-T Ethernet interfaces.

Port 0

- auto-negotiation enabled* speed 1000 or auto
- auto-negotiation enabled speed 100
- auto-negotiation disabled speed 100
- auto-negotiation disabled speed 100
- auto-negotiation enabled* speed 1000 or auto

Port 1

- auto-negotiation enabled* speed 1000 or auto
- auto-negotiation enabled speed 100
- auto-negotiation disabled speed 100
- auto-negotiation enabled speed 100
- auto-negotiation disabled speed 100

Link Status Between Port 1 and Port 2

- Up at 1000 Mb/s
- Up at 100 Mb/s
- Up at 100 Mb/s
- Down
- Down

* You cannot disable auto-negotiation when the speed is set to 1000 or auto.

Related Commands

speed (for 1000/10000/auto interfaces) — sets the link speed to 1000, 10000, or auto-negotiate the speed.

portmode hybrid

To accept *both* tagged and untagged frames, set a physical port or port-channel. A port configured this way is identified as a hybrid port in report displays.

Syntax

`portmode hybrid`

To return a port to accept *either* tagged or untagged frames (non-hybrid), use the `no portmode hybrid` command.

Defaults

non-hybrid

Command Modes

INTERFACE (conf-if-*interface-slot/port*)

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

The following describes the `interface` command shown in the following example. This example sets a port as hybrid, makes the port a tagged member of VLAN 20, and an untagged member of VLAN 10, which becomes the native VLAN of the port. The port now accepts:

- untagged frames and classify them as VLAN 10 frames
- VLAN 20 tagged frames

The following describes the `do show interfaces` command shown in the following example. This example shows output with "Hybrid" as the newly added value for 802.1QTagged. The options for this field are:

- True — port is tagged
- False — port is untagged
- Hybrid — port accepts both tagged and untagged frames

The following describes the `interface vlan` command shown in the following example. This example shows unconfiguration of the hybrid port using the `no portmode hybrid` command.

NOTE: Remove all other configurations on the port before you can remove the hybrid configuration from the port.

Example

```
Dell(conf)#interface tengig 0/20
Dell(conf-if-te-0/20)#no shut
Dell(conf-if-te-0/20)#portmode hybrid
Dell(conf-if-te-0/20)#sw
Dell(conf-if-te-0/20)#int vlan 10
Dell(conf-if-vl-10)#tag tengig 0/20
Dell(conf-if-vl-10)#int vlan 20
Dell(conf-if-vl-20)#untag tengig 0/20
Dell(conf-if-vl-20)#
```

Example (tagged hybrid)

```
Dell(conf)#interface tengig 0/20
Dell(conf-if-te-0/20)#no shut
Dell(conf-if-te-0/20)#portmode hybrid
Dell(conf-if-te-0/20#sw
Dell(conf-if-te-0/20)#int vlan 10
Dell(conf-if-vl-10)#int tengig 0/20
Dell(conf-if-vl-20)# untag tengig 0/20

Dell (conf-if-vl-20)#

Dell(conf)#do show interfaces switchport tengigabitethernet
3/20

Codes: U - Untagged, T - Tagged
       x - Dot1x untagged, X - Dot1x tagged
       G - GVRP tagged, M - Trunk, H - VSN tagged
       i - Internal untagged, I - Internal tagged, v - VLT
untagged,
```

```
       V - VLT tagged

Name: TenGigabitEthernet 3/20

802.1QTagged: Hybrid
Vlan membership:
Q   Vlans
U   20
T   10

Native VlanId: 20.

Dell(conf)#
```

Example (unconfigure the hybrid port)

```
Dell(conf-if-vl-20)#interface vlan 10
Dell(conf-if-vl-10)#no untagged tengig 0/20
Dell(conf-if-vl-10)#interface vlan 20
Dell(conf-if-vl-20)#no tagged tengig 0/20
Dell(conf-if-vl-20)#interface tengig 0/20
Dell(conf-if-te-0/20)#no portmode hybrid
Dell(conf-if-vl-20)#
```

Related Commands

show interfaces switchport — displays the configuration of switchport (Layer 2) interfaces on the switch.

vlan-stack trunk — specifies an interface as a trunk port to the Stackable VLAN network.

stack-unit portmode

Split a single 40G port into 4-10G ports on the MXL switch.

Syntax

```
stack-unit stack-unit port number portmode quad
```

Parameters

stack-unit Enter the stack member unit identifier of the stack member to reset. The range is 0 to 5.

NOTE: The MXL switch commands accept Unit ID numbers from 0 to 5, though the MXL switch supports stacking up to three units only with the Dell Networking OS version 8.3.7.1.

number Enter the port number of the 40G port to be split. Enter one of the following port numbers for the MXL switch: 48, 52, 56, or 60.

Defaults Disabled.

Command Modes CONFIGURATION

Command History

Version 9.2(0.0) Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

Splitting a 40G port into 4x10G port is supported on standalone and stacked units.

- You cannot use split ports as stack-link to stack an MXL Switch.
- The split ports MXL switch unit cannot be a part of any stacked system.
- The unit number with the split ports must be the default (stack-unit 0).
- This set up can be verified using `show system brief` command. If the unit ID is different than 0, it must be renumbered to 0 before ports are split by using the `stackunit id renumber 0` command in EXEC mode.

The quad port must be in a default configuration before it can be split into 4x10G ports. The 40G port is lost in the config when the port is split, so be sure that the port is also removed from other L2/L3 feature configurations.

The system must be reloaded after issuing the CLI for the change to take effect.

Port Channel Commands

A link aggregation group (LAG) is a group of links that appear to a MAC client as if they were a single link according to IEEE 802.3ad. In the Dell Networking OS, a LAG is referred to as a Port Channel.

- For the MXL switch, the maximum port channel ID is 128 and the maximum members per port channel is 16.

Because each port can be assigned to only one Port Channel, and each Port Channel must have at least one port, some of those nominally available Port Channels might have no function because they could have no members if there are not enough ports installed. In the MXL 10/40GbE Switch IO Module, those ports could be provided by stack members.

NOTE: The Dell Networking OS implementation of LAG or Port Channel requires that you configure a LAG on both switches manually. For information about Dell Networking OS link aggregation control protocol (LACP) for dynamic LAGs, refer to the *Link Aggregation Control Protocol (LACP)* chapter. For more information about configuring and using Port Channels, refer to the *Dell Networking OS Configuration Guide*.

channel-member

Add an interface to the Port Channel, while in INTERFACE PORTCHANNEL mode.

Syntax

`channel-member` *interface*

To delete an interface from a Port Channel, use the `no channel-member` *interface* command.

Parameters

interface	(OPTIONAL) Enter any of the following keywords and slot/port or number information: • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Defaults

Not configured.

Command Modes

INTERFACE PORTCHANNEL

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

Use the `interface port-channel` command to access this command.

You cannot add an interface to a Port Channel if the interface contains an IP address in its configuration.

Link MTU and IP MTU considerations for Port Channels are:

- All members must have the same link MTU value and the same IP MTU value.
- The Port Channel link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the channel members. For example, if the members have a link MTU of 2100 and an IP MTU 2000, the Port Channel's MTU values cannot be higher than 2100 for link MTU or 2000 bytes for IP MTU.

When an interface is removed from a Port Channel with the `no channel-member` command, the interface reverts to its configuration prior to joining the Port Channel.

An interface can belong to only one Port Channel.

You can add up to 16 interfaces to a Port Channel on the MXL switch. The interfaces can be located on different line cards but must be the same physical type and speed (for example, all 10-Gigabit Ethernet interfaces). However, you can combine 100/1000 interfaces and GE interfaces in the same Port Channel.

If the Port Channel contains a mix of interfaces with 100 Mb/s speed and 1000 Mb/s speed, the software disables those interfaces whose speed does not match the speed of the first interface configured and enabled in the Port Channel. If that first interface goes down, the Port Channel does not change its designated speed; disable and re-enable the Port Channel or change the order of the channel members configuration to change the designated speed. If the Port Channel contains a mix of interfaces with 100 Mb/s speed and 1000 Mb/s speed, the software disables those interfaces whose speed does not match the speed of the first interface configured and enabled in the Port Channel. If that first interface goes down, the Port Channel does not change its designated speed; disable and re-enable the Port Channel or change the order of the channel members configuration to change the designated speed. For more information about Port Channels, refer to the *Dell Networking OS Configuration Guide*.

Related Commands

description — assigns a descriptive text string to the interface.

interface port-channel — creates a Port Channel interface.

shutdown — disables/enables the port channel.

interface port-channel

Create a Port Channel interface, which is a link aggregation group (LAG) containing 16 physical interfaces on the XML switch.

Syntax

`interface port-channel channel-number`

To delete a Port Channel, use the `no interface port-channel channel-number` command.

Parameters

channel-number	For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128.

Defaults

Not configured.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

Port Channel interfaces are logical interfaces and can be either in Layer 2 mode (by using the `switchport` command) or Layer 3 mode (by configuring an IP address). You can add a Port Channel in Layer 2 mode to a VLAN.

A Port Channel can contain both 100/1000 interfaces and GE interfaces. Based on the first interface configured in the Port Channel and enabled, the Dell Networking OS determines if the Port Channel uses 100 Mb/s or 1000 Mb/s as the common speed. For more information, refer to channel-member.

If the line card is in a Jumbo mode chassis, you can also configure the `mtu` and `ip mtu` commands. The Link MTU and IP MTU values configured on the channel members must be greater than the Link MTU and IP MTU values configured on the Port Channel interface.

NOTE: In a Jumbo-enabled system, you must configure all members of a Port Channel with the same link MTU values and the same IP MTU values.

Example

```
Dell(conf)#int port-channel 2
Dell(conf-if-po-2)#
```

Related Commands

channel-member — adds a physical interface to the LAG.

interface — configures a physical interface.

interface loopback — configures a Loopback interface.

interface null — configures a null interface.

interface vlan — configures a VLAN.

shutdown — disables/enables the port channel.

minimum-links

Configure the minimum number of links in a LAG (Port Channel) that must be in "oper up" status for the LAG to be also in "oper up" status.

Syntax

```
minimum-links number
```

Parameters

number	Enter the number of links in a LAG that must be in "oper up" status. The range is from 1 to 16. The default is **1**.

Defaults

1

Command Modes

INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

If you use this command to configure the minimum number of links in a LAG that must be in "oper up" status, the LAG must have at least that number of "oper up" links before it can be declared as up. For example, if the required minimum is four, and only three are up, the LAG is considered down.

26

Internet Group Management Protocol (IGMP)

The IGMP commands are supported by the Dell Networking operating software (FTOS) on the Z9000 S4810 S4820T platform.

This chapter contains the following sections:

- IGMP Commands
- IGMP Snooping Commands

IGMP Commands

FTOS supports IGMPv1/v2/v3 and is compliant with RFC-3376.

Important Points to Remember

- FTOS supports protocol-independent multicast-sparse (PIM-SM) and protocol-independent source-specific multicast (PIM-SSM) include and exclude modes.
- IGMPv2 is the default version of IGMP on interfaces. You can configure IGMPv3 on interfaces. It is backward compatible with IGMPv2.
- On the S-Series, the maximum number of interfaces supported 31. On the S-Series, the maximum number of interfaces supported 31. The Z9000 supports up to 95 interfaces.
- There is no hard limit on the maximum number of groups supported.
- IGMPv3 router interoperability with IGMPv2 and IGMPv1 routers on the same subnet is not supported.
- An administrative command (`ip igmp version`) is added to manually set the IGMP version.
- All commands previously used for IGMPv2 are compatible with IGMPv3.

ip igmp group-join-limit

To limit the number of IGMP groups that can be joined in a second, use this feature.

Z9000 S4810 S4820TS6000

Syntax	`ip igmp group-join-limit` *number*	
Parameters	*number*	Enter the number of IGMP groups permitted to join in a second. The range is from 1 to 10000.
Defaults	none	
Command Modes	CONFIGURATION (conf-if-*interface-slot/port*)	

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.8.1.0	Introduced on the C-Series and S-Series.
Version 7.6.1.0	Introduced on the E-Series.

ip igmp last-member-query-interval

Change the last member query interval, which is the Max Response Time inserted into Group-Specific Queries sent in response to Leave Group messages. This interval is also the interval between Group-Specific Query messages.

Z9000 S4810 S4820TS6000

Syntax

`ip igmp last-member-query-interval` *`milliseconds`*

To return to the default value, use the `no ip igmp last-member-query-interval` command.

Parameters

milliseconds	Enter the number of milliseconds as the interval. The range is from 100 to 65535. The default is **1000 milliseconds**.

Defaults

1000 milliseconds

Command Modes

INTERFACE

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.

Version 7.8.1.0	Introduced on the S-Series.
Version 7.7.1.0	Introduced on the C-Series.
E-Series legacy command	

ip igmp querier-timeout

Change the interval that must pass before a multicast router decides that there is no longer another multicast router that should be the querier.

Z9000 S4810 S4820TS6000

Syntax

`ip igmp querier-timeout` *seconds*

To return to the default value, use the `no ip igmp querier-timeout` command.

Parameters

seconds	Enter the number of seconds the router must wait to become the new querier. The range is from 60 to 300. The default is **125 seconds**.

Defaults

125 seconds

Command Modes

INTERFACE

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.8.1.0	Introduced on the S-Series.
Version 7.7.1.0	Introduced on the C-Series.
Version 7.6.1.0	Introduced on the S-Series in Interface VLAN mode only to enable the system to act as an IGMP Proxy Querier.
Version 7.5.1.0	Introduced on the C-Series in Interface VLAN mode only to enable the system to act as an IGMP Proxy Querier.
E-Series legacy command.	

ip igmp query-interval

Change the transmission frequency of IGMP general queries the Querier sends.

Z9000 S4810 S4820TS6000

Syntax

`ip igmp query-interval` *seconds*

To return to the default values, use the `no ip igmp query-interval` command.

Parameters

seconds	Enter the number of seconds between queries sent out. The range is from 1 to 18000. The default is **60 seconds**.

Defaults

60 seconds

Command Modes

INTERFACE

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.8.1.0	Introduced on the S-Series.
Version 7.7.1.0	Introduced on the C-Series.
Version 7.6.1.0	Introduced on the S-Series in Interface VLAN mode only to enable the system to act as an IGMP Proxy Querier.
Version 7.5.1.0	Introduced on the C-Series in Interface VLAN mode only to enable the system to act as an IGMP Proxy Querier.
E-Series legacy command.	

ip igmp query-max-resp-time

Set the maximum query response time advertised in general queries.

Z9000 S4810 S4820TS6000

Syntax

`ip igmp query-max-resp-time` *seconds*

To return to the default values, use the `no ip igmp query-max-resp-time` command.

Parameters	*seconds*	Enter the number of seconds for the maximum response time. The range is from 1 to 25. The default is **10 seconds**.
Defaults	**10 seconds**	
Command Modes	INTERFACE	

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
Version 7.6.1.0	Introduced on the S-Series in Interface VLAN mode only to enable the system to act as an IGMP Proxy Querier.
Version 7.5.1.0	Introduced on the C-Series in Interface VLAN mode only to enable the system to act as an IGMP Proxy Querier.
E-Series legacy command.	

ip igmp version

Manually set the version of the router to IGMPv2 or IGMPv3.

Z9000 S4810 S4820TS6000

Syntax `ip igmp version {2 | 3}`

Parameters	**2**	Enter the number 2 to set the IGMP version number to IGMPv2.
	3	Enter the number 3 to set the IGMP version number to IGMPv3.
Defaults	**2** (that is, IGMPv2)	

Command Modes	INTERFACE
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*. The following is a list of the Dell Networking OS version history for this command.

Version 9.0.2.0	Introduced on the S6000.
Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.8.1.0	Introduced on the S-Series.
Version 7.7.1.0	Introduced on the C-Series.
Version 7.5.1.0	Introduced on the E-Series.

IGMP Snooping Commands

FTOS supports IGMP Snooping version 2 and 3 on all Dell Networking systems.

Important Points to Remember for IGMP Snooping

- FTOS supports version 1, version 2, and version 3 hosts.
- FTOS IGMP snooping implementation is based on IP multicast address (not based on Layer 2 multicast mac address) and the IGMP snooping entries are in Layer 3 flow table not in Layer 2 forwarding information base (FIB).
- FTOS IGMP snooping implementation is based on draft-ietf-magma-snoop-10.
- FTOS supports IGMP snooping on JUMBO-enabled cards.
- IGMP snooping is not enabled by default on the switch.
- A maximum of 1800 groups and 600 VLAN are supported.
- IGMP snooping is not supported on a default VLAN interface.
- IGMP snooping is not supported over VLAN-Stack-enabled VLAN interfaces (you must disable IGMP snooping on a VLAN interface before configuring VLAN-Stack-related commands).
- IGMP snooping does not react to Layer 2 topology changes triggered by spanning tree protocol (STP).
- IGMP snooping reacts to Layer 2 topology changes multiple spanning tree protocol (MSTP) triggers by sending a general query on the interface that comes in the FWD state.

Important Points to Remember for IGMP Querier

- The IGMP snooping Querier supports version 2.
- You must configure an IP address to the VLAN interface for IGMP snooping Querier to begin. The IGMP snooping Querier disables itself when a VLAN IP address is cleared, and then it restarts itself when an IP address is reassigned to the VLAN interface.

- When enabled, IGMP snooping Querier does not start if there is a statically configured multicast router interface in the VLAN.
- When enabled, IGMP snooping Querier starts after one query interval in case no IGMP general query (with IP SA lower than its VLAN IP address) is received on any of its VLAN members.
- When enabled, IGMP snooping Querier periodically sends general queries with an IP source address of the VLAN interface. If it receives a general query on any of its VLAN member, it checks the IP source address of the incoming frame.

  If the IP SA in the incoming IGMP general query frame is lower than the IP address of the VLAN interface, the switch disables its IGMP snooping Querier functionality.

  If the IP SA of the incoming IGMP general query is higher than the VLAN IP address, the switch continues to work as an IGMP snooping Querier.

ip igmp snooping enable

Enable IGMP snooping on all or a single VLAN. This command is the master on/off switch to enable IGMP snooping.

Z9000 S4810 S4820TS6000

Syntax	`ip igmp snooping enable` To disable IGMP snooping, use the `no ip igmp snooping enable` command.
Defaults	Disabled.
Command Modes	• CONFIGURATION • INTERFACE VLAN
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*. The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Usage Information	To enable IGMP snooping, enter this command. When you enable this command from CONFIGURATION mode, IGMP snooping enables on all VLAN interfaces (except the default VLAN).

NOTE: Execute the `no shutdown` command on the VLAN interface for IGMP Snooping to function.

Related Commands shutdown — (no shutdown) activates an interface.

ip igmp snooping fast-leave

Enable IGMP snooping fast-leave for this VLAN.

Z9000 S4810 S4820TS6000

Syntax `ip igmp snooping fast-leave`

To disable IGMP snooping fast leave, use the `no igmp snooping fast-leave` command.

Defaults Not configured.

Command Modes INTERFACE VLAN — (conf-if-vl-*n*)

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command.	

Usage Information Queriers normally send some queries when a leave message is received prior to deleting a group from the membership database. There may be situations when you require a fast deletion of a group. When you enable IGMP fast leave processing, the switch removes an interface from the multicast group as soon as it detects an IGMP version 2 leave message on the interface.

ip igmp snooping last-member-query-interval

The last member query interval is the maximum response time inserted into Group-Specific queries sent in response to Group-Leave messages.

Z9000 S4810 S4820TS6000

Syntax

`ip igmp snooping last-member-query-interval` *`milliseconds`*

To return to the default value, use the `no ip igmp snooping last-member-query-interval` command.

Parameters

milliseconds	Enter the interval in milliseconds. The range is from 100 to 65535. The default is **1000 milliseconds**.

Defaults

1000 milliseconds

Command Modes

INTERFACE VLAN

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Usage Information

This last-member-query-interval is also the interval between successive Group-Specific Query messages. To change the last-member-query interval, use this command.

ip igmp snooping mrouter

Statically configure a VLAN member port as a multicast router interface.

Z9000 S4810 S4820TS6000

Syntax

`ip igmp snooping mrouter interface` *`interface`*

To delete a specific multicast router interface, use the `no igmp snooping mrouter interface` *`interface`* command.

Parameters

interface *interface*	Enter the following keywords and slot/port or number information: • For a 100/1000 Ethernet interface, enter the keyword `gigabitethernet` followed by the slot/port information. • For a 1-Gigabit Ethernet interface, enter the keyword `gigabitethernet` followed by the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information. • For a Port Channel interface, enter the keywords `port-channel` then a number. For the C-Series and S-Series, the range is from 1 to 128.

Defaults Not configured.

Command Modes INTERFACE VLAN — (conf-if-vl-*n*)

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command.	

Usage Information Dell Networking OS provides the capability of statically configuring the interface to which a multicast router is attached. To configure a static connection to the multicast router, enter the `ip igmp snooping mrouter interface` command in the VLAN context. The interface to the router must be a part of the VLAN where you are entering the command.

ip igmp snooping querier

Enable IGMP querier processing for the VLAN interface.

Z9000 S4810 S4820TS6000

Syntax	`ip igmp snooping querier` To disable IGMP querier processing for the VLAN interface, use the `no ip igmp snooping querier` command.
Defaults	Not configured.
Command Modes	INTERFACE VLAN — (conf-if-vl-*n*)
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*. The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.2.0	Introduced on the S6000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Usage Information	This command enables the IGMP switch to send General Queries periodically. This behavior is useful when there is no multicast router present in the VLAN because the multicast traffic is not routed. Assign an IP address to the VLAN interface for the switch to act as a querier for this VLAN.

27

Layer 2

This chapter describes commands to configure Layer 2 features.
This chapter contains the following sections:

- MAC Addressing Commands
- Virtual LAN (VLAN) Commands

MAC Addressing Commands

The following commands are related to configuring, managing, and viewing MAC addresses.

mac-address-table aging-time

Specify an aging time for MAC addresses to remove from the MAC address table.

Syntax	`mac-address-table aging-time seconds`	
	To delete the configured aging time, use the `no mac-address-table aging-time seconds` command.	
Parameters	*seconds*	Enter either zero (0) or a number as the number of seconds before MAC addresses are relearned. To disable aging of the MAC address table, enter `0`. The range is from 10 to 1000000. The default is **1800 seconds**.
Defaults	**1800 seconds**	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Related Commands	mac learning-limit — sets the MAC address learning limits for a selected interface. show mac-address-table aging-time — displays the MAC aging time.	

mac-address-table static

Associate specific MAC or hardware addresses to an interface and virtual local area networks (VLANs).

Syntax

`mac-address-table static` *`mac-address`* `output` *`interface`* `vlan` *`vlan-id`*

To remove a MAC address, use the `no mac-address-table static` *`mac-address`* `output` *`interface`* `vlan` *`vlan-id`* command.

Parameters

mac-address	Enter the 48-bit hexadecimal address in nn:nn:nn:nn:nn:nn format.
output ***interface***	Enter the keyword `output` then one of the following interfaces for which traffic is forwarded: • For a Port Channel interface, enter the keywords `port-channel` then a number. The range is from 1 to 128. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
vlan ***vlan-id***	Enter the keyword `vlan` then a VLAN ID number from 1 to 4094.

Defaults Not configured.

Command Modes CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Related Commands show mac-address-table — displays the MAC address table.

mac-address-table station-move refresh-arp

Ensure that address resolution protocol (ARP) refreshes the egress interface when a station move occurs due to a topology change.

Syntax `[no] mac-address-table station-move refresh-arp`

Defaults none

Command Modes CONFIGURATION

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information For details about using this command, refer to the "NIC Teaming" section of the Layer 2 chapter in the *Dell Networking OS Configuration Guide*.

28
Link Aggregation Control Protocol (LACP)

This chapter contains commands for Dell Networks's implementation of the link aggregation control protocol (LACP) for creating dynamic link aggregation groups (LAGs) — known as "port-channels" in the Dell Networking operating software.

NOTE: For static LAG commands, refer to Port Channel Commands in the Interfaces chapter), based on the standards specified in the IEEE 802.3 Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.

lacp long-timeout

Configure a long timeout period (30 seconds) for an LACP session.

Syntax	`lacp long-timeout` To reset the timeout period to a short timeout (1 second), use the `no lacp long-timeout` command.	
Defaults	**1 second**	
Command Modes	INTERFACE (conf-if-po-number)	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	This command applies to dynamic port-channel interfaces only. When applied on a static port-channel, this command has no effect.	
Related Commands	show lacp — displays the LACP configuration.	

lacp port-priority

To influence which ports will be put in Standby mode when there is a hardware limitation that prevents all compatible ports from aggregating, configure the port priority.

Syntax	`lacp port-priority` *priority-value* To return to the default setting, use the `no lacp port-priority` *priority-value* command.	
Parameters	*priority-value*	Enter the port-priority value. The higher the value number, the lower the priority. The range is from 1 to 65535. The default is **32768**.

Defaults	**32768**	
Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

port-channel mode

Configure the LACP port channel mode.

Syntax `port-channel number mode [active] [passive] [off]`

Parameters

number	Enter the keywords `number` then a number.
active	Enter the keyword `active` to set the mode to the active state.
	NOTE: LACP modes are defined in *Usage Information*.
passive	Enter the keyword `passive` to set the mode to the passive state.
	NOTE: LACP modes are defined in *Usage Information*.
off	Enter the keyword `off` to set the mode to the off state.
	NOTE: LACP modes are defined in *Usage Information*.

Defaults	**off**	
Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

LACP Modes

Mode	**Function**
active	An interface is in an active negotiating state in this mode. LACP runs on any link configured in the active state and also automatically initiates negotiation with other ports by initiating LACP packets.

Mode	Function
passive	An interface is not in an active negotiating state in this mode. LACP runs on any link configured in the passive state. Ports in a passive state respond to negotiation requests from other ports that are in active states. Ports in a passive state respond to LACP packets
off	An interface cannot be part of a dynamic port channel in off mode. LACP does not run on a port configured in off mode.

port-channel-protocol lacp

Enable LACP on any LAN port.

Syntax

```
port-channel-protocol lacp
```

To disable LACP on a LAN port, use the `no port-channel-protocol lacp` command.

Command Modes

INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell(conf)#interface TenGigabitethernet 3/15
Dell(conf-if-tengig-3/15)#no shutdown
Dell(conf-if-tengig-3/15)#port-channel-protocol lacp
Dell(conf-if-tengig-3/15-lacp)#port-channel 32 mode active
...
Dell(conf)#interface TenGigabitethernet 3/16
Dell(conf-if-tengig-3/16)#no shutdown
Dell(conf-if-tengig-3/16)#port-channel-protocol lacp
Dell(conf-if-tengig-3/16-lacp)#port-channel 32 mode active
```

Related Commands

show lacp — displays the LACP information.

show interfaces port-channel — displays information on configured Port Channel groups.

Configuration Tasks for Port Channel Interfaces

To configure a port channel (LAG), use the commands similar to those found in physical interfaces. By default, no port channels are configured in the startup configuration.

These are the mandatory and optional configuration tasks:

- Creating a Port Channel (mandatory)
- Adding a Physical Interface to a Port Channel (mandatory)
- Reassigning an Interface to a New Port Channel (optional)

- Configuring the Minimum Oper Up Links in a Port Channel (optional)
- Adding or Removing a Port Channel from a VLAN (optional)
- Assigning an IP Address to a Port Channel (optional)
- Deleting or Disabling a Port Channel (optional)
- Load Balancing Through Port Channels (optional)

Creating a Port Channel

You can create up to 128 port channels with eight port members per group on the Z9000 S4810 S4820T.

To configure a port channel, use the following commands.

1. Create a port channel.
   CONFIGURATION mode
   `interface port-channel id-number`
2. Ensure that the port channel is active.
   INTERFACE PORT-CHANNEL mode
   `no shutdown`

After you enable the port channel, you can place it in Layer 2 or Layer 3 mode. To place the port channel in Layer 2 mode or configure an IP address to place the port channel in Layer 3 mode, use the `switchport` command.

You can configure a port channel as you would a physical interface by enabling or configuring protocols or assigning access control lists.

Adding a Physical Interface to a Port Channel

The physical interfaces in a port channel can be on any line card in the chassis, but must be the same physical type.

NOTE: Port channels can contain a mix of Gigabit Ethernet and 10/100/1000 Ethernet interfaces, but FTOS disables the interfaces that are not the same speed of the first channel member in the port channel (refer to 10/100/1000 Mbps Interfaces in Port Channels).

You can add any physical interface to a port channel if the interface configuration is minimal. You can configure only the following commands on an interface if it is a member of a port channel:

- `description`
- `shutdown/no shutdown`
- `mtu`
- `ip mtu` (if the interface is on a Jumbo-enabled by default)

NOTE:

A logical port channel interface cannot have flow control. Flow control can only be present on the physical interfaces if they are part of a port channel.

NOTE: The S4810 S4820T supports jumbo frames by default (the default maximum transmission unit (MTU) is 1554 bytes). The Z9000 supports jumbo frames by default (the default maximum transmission unit (MTU) is 12000 bytes). To configure the MTU, use the `mtu` command from INTERFACE mode.

To view the interface's configuration, enter INTERFACE mode for that interface and use the `show config` command or from EXEC Privilege mode, use the `show running-config interface` *`interface`* command.

When an interface is added to a port channel, FTOS recalculates the hash algorithm.

To add a physical interface to a port, use the following commands.

1. Add the interface to a port channel.
   INTERFACE PORT-CHANNEL mode
   `channel-member` *`interface`*
   The *`interface`* variable is the physical interface type and slot/port information.
2. Double check that the interface was added to the port channel.
   INTERFACE PORT-CHANNEL mode
   `show config`

To view the port channel's status and channel members in a tabular format, use the `show interfaces port-channel brief` command in EXEC Privilege mode, as shown in the following example.

Example of the `show interfaces port-channel brief` Command

```
FTOS#show int port brief

LAG Mode Status Uptime   Ports
1 L2L3   up     00:06:03 Gi 13/6 (Up) *
                         Gi 13/12 (Up)
2 L2L3   up     00:06:03 Gi 13/7 (Up) *
                         Gi 13/8 (Up)
                         Gi 13/13 (Up)
                         Gi 13/14 (Up)
FTOS#
```

The following example shows the port channel's mode (L2 for Layer 2 and L3 for Layer 3 and L2L3 for a Layer 2-port channel assigned to a routed VLAN), the status, and the number of interfaces belonging to the port channel.

Example of the `show interface port-channel` Command

```
FTOS>show interface port-channel 20
Port-channel 20 is up, line protocol is up
Hardware address is 00:01:e8:01:46:fa
Internet address is 1.1.120.1/24
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 2000 Mbit
Members in this channel: Gi 9/10 Gi 9/17
ARP type: ARPA, ARP timeout 04:00:00
Last clearing of "show interface" counters 00:00:00
Queueing strategy: fifo
  1212627 packets input, 1539872850 bytes
  Input 1212448 IP Packets, 0 Vlans 0 MPLS
  4857 64-byte pkts, 17570 over 64-byte pkts, 35209 over 127-byte pkts
  69164 over 255-byte pkts, 143346 over 511-byte pkts, 942523 over 1023-byte
pkts
  Received 0 input symbol errors, 0 runts, 0 giants, 0 throttles
  42 CRC, 0 IP Checksum, 0 overrun, 0 discarded
  2456590833 packets output, 203958235255 bytes, 0 underruns
  Output 1640 Multicasts, 56612 Broadcasts, 2456532581 Unicasts
  2456590654 IP Packets, 0 Vlans, 0 MPLS
  0 throttles, 0 discarded
Rate info (interval 5 minutes):
  Input 00.01Mbits/sec, 2 packets/sec
```

```
  Output 81.60Mbits/sec, 133658 packets/sec
Time since last interface status change: 04:31:57

FTOS>
```

When more than one interface is added to a Layer 2-port channel, FTOS selects one of the active interfaces in the port channel to be the primary port. The primary port replies to flooding and sends protocol data units (PDUs). An asterisk in the `show interfaces port-channel brief` command indicates the primary port.

As soon as a physical interface is added to a port channel, the properties of the port channel determine the properties of the physical interface. The configuration and status of the port channel are also applied to the physical interfaces within the port channel. For example, if the port channel is in Layer 2 mode, you cannot add an IP address or a static MAC address to an interface that is part of that port channel. In the following example, interface GigabitEthernet 1/6 is part of port channel 5, which is in Layer 2 mode, and an error message appeared when an IP address was configured.

Example of Error Due to an Attempt to Configure an Interface that is Part of a Port Channel

```
FTOS(conf-if-portch)#show config
!
interface Port-channel 5
  no ip address
  switchport
  channel-member GigabitEthernet 1/6
FTOS(conf-if-portch)#int gi 1/6
FTOS(conf-if)#ip address 10.56.4.4 /24
% Error: Port is part of a LAG Gi 1/6.
FTOS(conf-if)#
```

Reassigning an Interface to a New Port Channel

An interface can be a member of only one port channel. If the interface is a member of a port channel, remove it from the first port channel and then add it to the second port channel.

Each time you add or remove a channel member from a port channel, FTOS recalculates the hash algorithm for the port channel.

To reassign an interface to a new port channel, use the following commands.

1. Remove the interface from the first port channel.
   INTERFACE PORT-CHANNEL mode
   `no channel-member` *`interface`*
2. Change to the second port channel INTERFACE mode.
   INTERFACE PORT-CHANNEL mode
   `interface port-channel id` *`number`*
3. Add the interface to the second port channel.
   INTERFACE PORT-CHANNEL mode
   `channel-member interface`

The following example shows moving the GigabitEthernet 1/8 interface from port channel 4 to port channel 3.

Example of Moving an Interface to a New Port Channel

```
FTOS(conf-if-portch)#show config
!
interface Port-channel 4
```

```
  no ip address
  channel-member GigabitEthernet 1/8
  no shutdown
FTOS(conf-if-portch)#no chann gi 1/8
FTOS(conf-if-portch)#int port 5
FTOS(conf-if-portch)#channel gi 1/8
FTOS(conf-if-portch)#sho conf
!
interface Port-channel 5
  no ip address
  channel-member GigabitEthernet 1/8
  shutdown
FTOS(conf-if-portch)#
```

Configuring the Minimum Oper Up Links in a Port Channel

You can configure the minimum links in a port channel (LAG) that must be in "oper up" status to consider the port channel to be in "oper up" status.

To set the "oper up" status of your links, use the following command.

- Enter the number of links in a LAG that must be in "oper up" status.

  INTERFACE mode

  `minimum-links` *number*

  The default is **1**.

Example of Configuring the Minimum Oper Up Links in a Port Channel

```
FTOS#config t
FTOS(conf)#int po 1
FTOS(conf-if-po-1)#minimum-links 5
FTOS(conf-if-po-1)#
```

Adding or Removing a Port Channel from a VLAN

As with other interfaces, you can add Layer 2 port channel interfaces to VLANs. To add a port channel to a VLAN, place the port channel in Layer 2 mode (by using the `switchport` command).

To add or remove a VLAN port channel and to view VLAN port channel members, use the following commands.

- Add the port channel to the VLAN as a tagged interface.

  INTERFACE VLAN mode

  `tagged port-channel` *id number*

  An interface with tagging enabled can belong to multiple VLANs.
- Add the port channel to the VLAN as an untagged interface.

  INTERFACE VLAN mode

  `untagged port-channel` *id number*

  An interface without tagging enabled can belong to only one VLAN.
- Remove the port channel with tagging enabled from the VLAN.

  INTERFACE VLAN mode

  `no tagged port-channel` *id number*

  or

```
no untagged port-channel id number
```

- Identify which port channels are members of VLANs.

  EXEC Privilege mode

  ```
  show vlan
  ```

Configuring VLAN Tags for Member Interfaces

To configure and verify VLAN tags for individual members of a port channel, perform the following:

1. Configure VLAN membership on individual ports

   INTERFACE mode

   ```
   FTOS(conf-if-te-0/2)#vlan tagged 2,3-4
   ```

2. Use the `switchport` command in INTERFACE mode to enable Layer 2 data transmissions through an individual interface

   INTERFACE mode

   ```
   FTOS(conf-if-te-0/2)#switchport
   ```

3. Verify the manually configured VLAN membership (**show interfaces switchport** *interface* command).

   EXEC mode

   ```
   FTOS(conf)# interface tengigabitethernet 0/1
   FTOS(conf-if-te-0/1)#switchport
   FTOS(conf-if-te-0/1)# vlan tagged 2-5,100,4010
   FTOS#show interfaces switchport te 0/1

   Codes:  U - Untagged, T - Tagged
           x - Dot1x untagged, X - Dot1x tagged
           G - GVRP tagged, M - Trunk, H - VSN tagged
           i - Internal untagged, I - Internal tagged, v - VLT untagged, V -
   VLT tagged

   Name: TenGigabitEthernet 0/1
   802.1QTagged: True
   Vlan membership:
   Q       Vlans
   T       2-5,100,4010

   FTOS#
   ```

Deleting or Disabling a Port Channel

To delete or disable a port channel, use the following commands.

- Delete a port channel.

  CONFIGURATION mode

  ```
  no interface portchannel channel-number
  ```

- Disable a port channel.

  ```
  shutdown
  ```

  When you disable a port channel, all interfaces within the port channel are operationally down also.

29

Link Layer Discovery Protocol (LLDP)

Link layer discovery protocol (LLDP) advertises connectivity and management from the local station to the adjacent stations on an IEEE 802 LAN.

LLDP facilitates multi-vendor interoperability by using standard management tools to discover and make available a physical topology for network management. The Dell Networking operating software implementation of LLDP is based on IEEE standard 801.1ab.

The starting point for using LLDP is invoking LLDP with the `protocol lldp` command in either CONFIGURATION or INTERFACE mode.

The information LLDP distributes is stored by its recipients in a standard management information base (MIB). You can access the information by a network management system through a management protocol such as simple network management protocol (SNMP).

For details about implementing LLDP/LLDP-MED, refer to the Link Layer Discovery Protocol chapter of the *Dell Networking OS Configuration Guide*.

advertise dot1-tlv

Advertise dot1 TLVs (Type, Length, Value).

Syntax

```
advertise dot1-tlv {port-protocol-vlan-id | port-vlan-id | vlan-name}
```

To remove advertised dot1-tlv, use the `no advertise dot1-tlv {port-protocol-vlan-id | port-vlan-id | vlan-name}` command.

Parameters

port-protocol-vlan-id	Enter the keywords `port-protocol-vlan-id` to advertise the port protocol VLAN identification TLV.
port-vlan-id	Enter the keywords `port-vlan-id` to advertise the port VLAN identification TLV.
vlan-name	Enter the keywords `vlan-name` to advertise the vlan-name TLV.

Defaults Disabled.

Command Modes CONFIGURATION (conf-lldp) and INTERFACE (conf-if-*interface*-lldp)

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Related Commands	protocol lldp (Configuration) — enables LLDP globally. debug lldp interface — debugs LLDP. show lldp neighbors — displays the LLDP neighbors. show running-config lldp — displays the LLDP running configuration.

advertise dot3-tlv

Advertise dot3 TLVs (Type, Length, Value).

Syntax	`advertise dot3-tlv {max-frame-size}` To remove advertised dot3-tlv, use the `no advertise dot3-tlv {max-frame-size}` command.
Parameters	**max-frame-size** — Enter the keywords `max-frame-size` to advertise the dot3 maximum frame size.
Defaults	none
Command Modes	CONFIGURATION (conf-lldp) and INTERFACE (conf-if-*interface*-lldp)
Command History	**Version 9.2(0.0)** — Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only. **Version 8.3.16.1** — Introduced on the MXL 10/40GbE Switch IO Module.

advertise management-tlv

Advertise management TLVs (Type, Length, Value).

Syntax	`advertise management-tlv {system-capabilities | system-description | system-name}` To remove advertised management TLVs, use the `no advertise management-tlv {system-capabilities | system-description | system-name}` command.
Parameters	**system-capabilities** — Enter the keywords `system-capabilities` to advertise the system capabilities TLVs to the LLDP peer. **system-description** — Enter the keywords `system-description` to advertise the system description TLVs to the LLDP peer. **system-name** — Enter the keywords `system-name` to advertise the system name TLVs to the LLDP peer.
Defaults	none

Command Modes	CONFIGURATION (conf-lldp)	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	The command options `system-capabilities`, `system-description`, and `system-name` can be invoked individually or together, in any sequence.	

clear lldp counters

Clear LLDP transmitting and receiving counters for all physical interfaces or a specific physical interface.

Syntax `clear lldp counters` *interface*

Parameters

interface	Enter the following keywords and slot/port or number information: • For a 10-Gigabit Ethernet interface, enter the keyword `tenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Defaults	none	
Command Modes	EXEC Privilege	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

clear lldp neighbors

Clear LLDP neighbor information for all interfaces or a specific interface.

Syntax `clear lldp neighbors {`*interface*`}`

Parameters

interface	Enter the following keywords and slot/port or number information: • For a 10-Gigabit Ethernet interface, enter the keyword `tenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Defaults	none	
Command Modes	EXEC Privilege	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

debug lldp interface

To display timer events, neighbor additions or deletions, and other information about incoming and outgoing packets, enable LLDP debugging.

Syntax

`debug lldp interface {`*`interface`*` | all}{events | packet {brief | detail} {tx | rx | both}}`

To disable debugging, use the `no debug lldp interface {`*`interface`*` | all}{events} {packet {brief | detail} {tx | rx | both}}` command.

Parameters

interface	Enter the following keywords and slot/port or number information: • For a 10-Gigabit Ethernet interface, enter the keyword `tenGigabitEthernet` then the slot/port information. • For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
all	(OPTIONAL) Enter the keyword `all` to display information on all interfaces.
events	(OPTIONAL) Enter the keyword `events` to display major events such as timer events.
packet	(OPTIONAL) Enter the keyword `packet` to display information regarding packets coming in or going out.
brief	(OPTIONAL) Enter the keyword `brief` to display brief packet information.
detail	(OPTIONAL) Enter the keyword `detail` to display detailed packet information.
tx	(OPTIONAL) Enter the keyword `tx` to display transmit-only packet information.
rx	(OPTIONAL) Enter the keyword `rx` to display receive-only packet information.
both	(OPTIONAL) Enter the keyword `both` to display both receive and transmit packet information.

Defaults	none	
Command Modes	EXEC Privilege	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

disable

Enable or disable LLDP.

Syntax	`disable` To enable LLDP, use the `no disable` command.	
Defaults	Enabled, that is `no disable`.	
Command Modes	CONFIGURATION (conf-lldp) and INTERFACE (conf-if-*interface*-lldp)	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Related Commands	protocol lldp (Configuration) — enables LLDP globally. debug lldp interface — debugs LLDP. show lldp neighbors — displays the LLDP neighbors. show running-config lldp — displays the LLDP running configuration.	

hello

Configure the rate at which the LLDP control packets are sent to its peer.

Syntax	`hello` *`seconds`* To revert to the default, use the `no hello` *`seconds`* command.	
Parameters	***seconds***	Enter the rate, in seconds, at which the control packets are sent to its peer. The rate is from 5 to 180 seconds. The default is **30 seconds**.
Defaults	**30 seconds**	
Command Modes	CONFIGURATION (conf-lldp) and INTERFACE (conf-if-*interface*-lldp)	

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

mode

To receive or transmit, set LLDP.

Syntax	`mode {tx | rx}` To return to the default, use the `no mode {tx | rx}` command.	
Parameters	**tx**	Enter the keyword `tx` to set the mode to transmit.
	rx	Enter the keyword `rx` to set the mode to receive.
Defaults	Both **transmit** and **receive**.	
Command Modes	CONFIGURATION (conf-lldp) and INTERFACE (conf-if-*interface*-lldp)	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Related Commands	protocol lldp (Configuration) — enables LLDP globally. show lldp neighbors — displays the LLDP neighbors.	

multiplier

Set the number of consecutive misses before LLDP declares the interface dead.

Syntax	`multiplier integer` To return to the default, use the `no multiplier integer` command.	
Parameters	***integer***	Enter the number of consecutive misses before the LLDP declares the interface dead. The range is from 2 to 10.
Defaults	**4 x hello**	
Command Modes	CONFIGURATION (conf-lldp) and INTERFACE (conf-if-*interface*-lldp)	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Configure LLDP

Configuring LLDP is a two-step process.

1. Enable LLDP globally.
2. Advertise TLVs out of an interface.

Related Configuration Tasks

- Viewing the LLDP Configuration
- Viewing Information Advertised by Adjacent LLDP Agents
- Configuring LLDPDU Intervals
- Configuring Transmit and Receive Mode
- Configuring a Time to Live
- Debugging LLDP

Important Points to Remember

- LLDP is enabled by default.
- Dell Networking systems support up to eight neighbors per interface.
- Dell Networking systems support a maximum of 8000 total neighbors per system. If the number of interfaces multiplied by eight exceeds the maximum, the system does not configure more than 8000.
- INTERFACE level configurations override all CONFIGURATION level configurations.
- LLDP is not hitless.

LLDP Compatibility

- Spanning tree and force10 ring protocol "blocked" ports allow LLDPDUs.
- 802.1X controlled ports do not allow LLDPDUs until the connected device is authenticated.

CONFIGURATION versus INTERFACE Configurations

All LLDP configuration commands are available in PROTOCOL LLDP mode, which is a sub-mode of the CONFIGURATION mode and INTERFACE mode.

- Configurations made at the CONFIGURATION level are global; that is, they affect all interfaces on the system.
- Configurations made at the INTERFACE level affect only the specific interface; they override CONFIGURATION level configurations.

Example of the `protocol lldp` Command (CONFIGURATION Level)

```
R1(conf)#protocol lldp
R1(conf-lldp)#?
advertise        Advertise TLVs
disable          Disable LLDP protocol globally
end              Exit from configuration mode
exit             Exit from LLDP configuration mode
hello            LLDP hello configuration
mode             LLDP mode configuration (default = rx and tx)
```

```
multiplier       LLDP multiplier configuration
no               Negate a command or set its defaults
show             Show LLDP configuration


R1(conf-lldp)#exit
R1(conf)#interface gigabitethernet 1/31
R1(conf-if-gi-1/31)#protocol lldp
R1(conf-if-gi-1/31-lldp)#?
advertise        Advertise TLVs
disable          Disable LLDP protocol on this interface
end              Exit from configuration mode
exit             Exit from LLDP configuration mode
hello            LLDP hello configuration
mode             LLDP mode configuration (default = rx and tx)
multiplier       LLDP multiplier configuration
no               Negate a command or set its defaults
show             Show LLDP configuration
R1(conf-if-gi-1/31-lldp)#
```

Enabling LLDP

LLDP is enabled by default. Enable and disable LLDP globally or per interface. If you enable LLDP globally, all UP interfaces send periodic LLDPDUs.

To enable LLDP, use the following command.

1. Enter Protocol LLDP mode.

   CONFIGURATION or INTERFACE mode

   `protocol lldp`
2. Enable LLDP.

   PROTOCOL LLDP mode

   `no disable`

Disabling and Undoing LLDP

To disable or undo LLDP, use the following command.

- Disable LLDP globally or for an interface.

  `disable`

To undo an LLDP configuration, precede the relevant command with the keyword `no`.

Enabling LLDP on Management Ports

LLDP on management ports is enabled by default.

To enable LLDP on management ports, use the following command.

1. Enter Protocol LLDP mode.
   CONFIGURATION mode
   ```
   protocol lldp
   ```
2. Enable LLDP.
   PROTOCOL LLDP mode
   ```
   no disable
   ```

Disabling and Undoing LLDP on Management Ports

To disable or undo LLDP on management ports, use the following command.

1. Enter Protocol LLDP mode.
   CONFIGURATION mode.
   ```
   protocol lldp
   ```
2. Enter LLDP management-interface mode.
   LLDP-MANAGEMENT-INTERFACE mode.
   ```
   management-interface
   ```
3. Enter the `disable` command.
   LLDP-MANAGEMENT-INTERFACE mode.

To undo an LLDP management port configuration, precede the relevant command with the keyword `no`.

Advertising TLVs

You can configure the system to advertise TLVs out of all interfaces or out of specific interfaces.

- If you configure the system globally, all interfaces send LLDPDUs with the specified TLVs.
- If you configure an interface, only the interface sends LLDPDUs with the specified TLVs.
- If you configure LLDP both globally and at interface level, the interface level configuration overrides the global configuration.

To advertise TLVs, use the following commands.

1. Enter LLDP mode.
   CONFIGURATION or INTERFACE mode
   `protocol lldp`
2. Advertise one or more TLVs.
   PROTOCOL LLDP mode
   `advertise {management-tlv | dot1-tlv | dot3-tlv | med}`
   Include the keyword for each TLV you want to advertise.
   - For management TLVs: `system-capabilities, system-description`.
   - For 802.1 TLVs: `port-protocol-vlan-id, port-vlan-id vlan-name`.
   - For 802.3 TLVs: `max-frame-size`.
   - For TIA-1057 TLVs:
     * `guest-voice`
     * `guest-voice-signaling`
     * `location-identification`
     * `power-via-mdi`
     * `softphone-voice`
     * `streaming-video`
     * `video-conferencing`
     * `video-signaling`
     * `voice`
     * `voice-signaling`

   NOTE: The `vlan-name` command is supported on S4810 S4820T systems only.

In the following example, LLDP is enabled globally. R1 and R2 are transmitting periodic LLDPDUs that contain management, 802.1, and 802.3 TLVs.

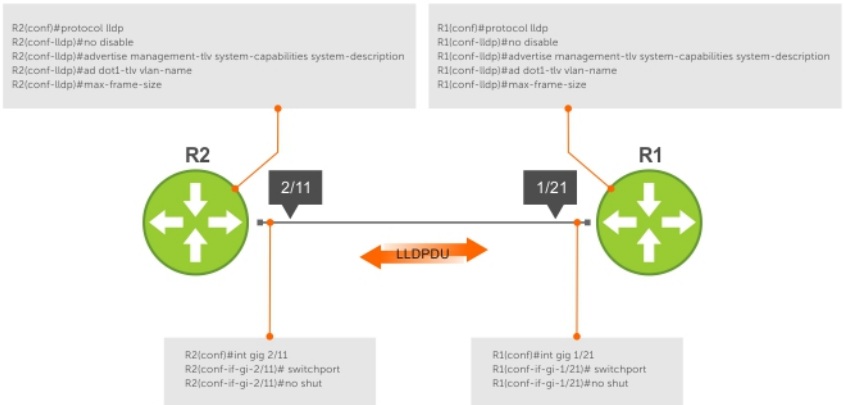

Figure 2. Configuring LLDP

Viewing the LLDP Configuration

To view the LLDP configuration, use the following command.

- Display the LLDP configuration.
  CONFIGURATION or INTERFACE mode
  ```
  show config
  ```

Example of Viewing LLDP Global Configurations

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  hello 10
  no disable
R1(conf-lldp)#
```

Example of Viewing LLDP Interface Configurations

```
R1(conf-lldp)#exit
R1(conf)#interface gigabitethernet 1/31
R1(conf-if-gi-1/31)#show config
!
interface GigabitEthernet 1/31
  no ip address
  switchport
  no shutdown
R1(conf-if-gi-1/31)#protocol lldp
R1(conf-if-gi-1/31-lldp)#show config
!
  protocol lldp
R1(conf-if-gi-1/31-lldp)#
```

Viewing Information Advertised by Adjacent LLDP Agents

To view brief information about adjacent devices or to view all the information that neighbors are advertising, use the following commands.

- Display brief information about adjacent devices.
  ```
  show lldp neighbors
  ```
- Display all of the information that neighbors are advertising.
  ```
  show lldp neighbors detail
  ```

Example of Viewing Brief Information Advertised by Neighbors

```
R1(conf-if-gi-1/31-lldp)#end
R1(conf-if-gi-1/31)#do show lldp neighbors
Loc PortID   Rem Host Name  Rem Port Id            Rem Chassis Id
-------------------------------------------------------------------------
Gi 1/21      -              GigabitEthernet 2/11   00:01:e8:06:95:3e
Gi 1/31      -              GigabitEthernet 3/11   00:01:e8:09:c2:4a
```

Example of Viewing Details Advertised by Neighbors

```
R1#show lldp neighbors detail
============================================================================
Local Interface Gi 1/21 has 1 neighbor
  Total Frames Out: 6547
  Total Frames In: 4136
  Total Neighbor information Age outs: 0
  Total Frames Discarded: 0
  Total In Error Frames: 0
  Total Unrecognized TLVs: 0
  Total TLVs Discarded: 0
  Next packet will be sent after 7 seconds
  The neighbors are given below:
----------------------------------------------------------------------------

  Remote Chassis ID Subtype: Mac address (4)
  Remote Chassis ID: 00:01:e8:06:95:3e
  Remote Port Subtype: Interface name (5)
  Remote Port ID: GigabitEthernet 2/11
  Local Port ID: GigabitEthernet 1/21
  Locally assigned remote Neighbor Index: 4
  Remote TTL: 120
  Information valid for next 120 seconds
  Time since last information change of this neighbor: 01:50:16
  Remote MTU: 1554
  Remote System Desc: Dell Force10 Networks Real Time Operating System Software
    . Dell Force10 Operating System Version: 1.0. Dell Force10 App
      lication Software Version: 7.5.1.0. Copyright (c) 19
      99-Build Time: Thu Aug 9 01:05:51 PDT 2007
  Existing System Capabilities: Repeater Bridge Router
  Enabled System Capabilities: Repeater Bridge Router
  Remote Port Vlan ID: 1
  Port and Protocol Vlan ID: 1, Capability: Supported, Status: Enabled
  ------------------------------------------------------------------------

============================================================================
```

Configuring LLDPDU Intervals

LLDPDUs are transmitted periodically; the default interval is **30 seconds**.

To configure LLDPDU intervals, use the following command.

- Configure a non-default transmit interval.

  CONFIGURATION mode or INTERFACE mode

  `hello`

Example of Viewing LLDPDU Intervals

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#mode ?
rx          Rx only
tx          Tx only
R1(conf-lldp)#mode tx
```

```
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  mode tx
  no disable
R1(conf-lldp)#no mode
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#
```

Configuring Transmit and Receive Mode

After you enable LLDP, Dell Networking systems transmit *and* receive LLDPDUs by default.

To configure the system to transmit or receive only and return to the default, use the following commands.

- Transmit only.

  CONFIGURATION mode or INTERFACE mode

  ```
  mode tx
  ```
- Receive only.

  CONFIGURATION mode or INTERFACE mode

  ```
  mode rx
  ```
- Return to the default setting.

  CONFIGURATION mode or INTERFACE mode

  ```
  no mode
  ```

Example of Configuring a Single Mode

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#mode ?
rx              Rx only
tx              Tx only
R1(conf-lldp)#mode tx
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  mode tx
  no disable
R1(conf-lldp)#no mode
R1(conf-lldp)#show config
```

```
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#
```

Configuring a Time to Live

The information received from a neighbor expires after a specific amount of time (measured in seconds) called a time to live (TTL).

The TTL is the product of the LLDPDU transmit interval (hello) and an integer called a multiplier. The default multiplier is **4**, which results in a default TTL of 120 seconds.

- Adjust the TTL value.

  CONFIGURATION mode or INTERFACE mode.

  ```
  multiplier
  ```
- Return to the default multiplier value.

  CONFIGURATION mode or INTERFACE mode.

  ```
  no multiplier
  ```

Example of the `multiplier` Command to Configure Time to Live

```
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#multiplier ?
<2-10>            Multiplier (default=4)
R1(conf-lldp)#multiplier 5
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  multiplier 5
  no disable
R1(conf-lldp)#no multiplier
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#
```

30

Quality of Service (QoS)

The Dell Networking operating software (FTOS) commands for quality of service (QoS) include traffic conditioning and congestion control. QoS commands are supported on the I/O Aggregator Z-Series S4810 S4820T platform.

This chapter contains the following sections:

- Global Configuration Commands
- Per-Port QoS Commands
- Policy-Based QoS Commands

Per-Port QoS Commands

Per-port QoS (port-based QoS) allows you to define the QoS configuration on a per-physical-port basis.

dot1p-priority

Assign a value to the IEEE 802.1p bits on the traffic this interface receives.

Syntax

`dot1p-priority` *priority-value*

To delete the IEEE 802.1p configuration on the interface, use the `no dot1p-priority` command.

Parameters

priority-value Enter a value from 0 to 7.

dot1p	Queue Number
0	2
1	0
2	1
3	3
4	4
5	5
6	6
7	7

Defaults none

Command Modes INTERFACE

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

The `dot1p-priority` command changes the priority of incoming traffic on the interface. The system places traffic marked with a priority in the correct queue and processes that traffic according to its queue.

When you set the priority for a port channel, the physical interfaces assigned to the port channel are configured with the same value. You cannot assign the `dot1p-priority` command to individual interfaces in a port channel.

rate shape

Shape the traffic output on the selected interface.

Syntax

```
rate shape [kbps] rate [burst-KB]
```

Parameters

kbps	Enter the keyword `kbps` to specify the rate limit in Kilobits per second (Kbps). Make the following value a multiple of 64. The range is from 0 to 40000000. The default granularity is Megabits per second (Mbps).
rate	Enter the outgoing rate in multiples of 10 Mbps. The range is from 10 to 10000.
burst-KB	(OPTIONAL) Enter the burst size in KB. The range is from 0 to 10000. The default is **50**.

Defaults Granularity for rate is **Mbps** unless you use the `kbps` option.

Command Modes INTERFACE

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Related Commands rate-shape — shapes traffic output as part of the designated policy.

service-class dynamic dot1p

Honor all 802.1p markings on incoming switched traffic on an interface (from INTERFACE mode) or on all interfaces (from CONFIGURATION mode). A CONFIGURATION mode entry supersedes an INTERFACE mode entry.

Syntax

```
service-class dynamic dot1p
```

To return to the default setting, use the `no service-class dynamic dot1p` command.

Defaults

All dot1p traffic is mapped to Queue 0 unless you enable the `service-class dynamic dot1p` command. The default mapping is as follows:

dot1p	Queue ID
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

Command Modes

- INTERFACE
- CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To honor all incoming 802.1p markings on incoming switched traffic on the interface, enter this command. By default, this facility is not enabled (that is, the 802.1p markings on incoming traffic are not honored).

You can apply this command on both physical interfaces and port channels. When you set the service-class dynamic for a port channel, the physical interfaces assigned to the port channel are automatically configured; you cannot assign the `service-class dynamic` command to individual interfaces in a port channel.

- All dot1p traffic is mapped to Queue 0 unless you enable the `service-class dynamic dot1p` command on an interface or globally.
- Layer 2 or Layer 3 service policies supersede dot1p service classes.

service-class dot1p-mapping

Configure a service-class criterion based on a dot1p value.

Z9000 S4810 S4820T

Syntax

```
service-class dot1p-mapping {dot1p0 value | dot1p1 value |
dot1p2 value | dot1p3 value | dot1p4 value| dot1p5 value |
dot1p6 value | dot1p7 value}
```

Parameters

dot1p0 *value* ... dot1p7 *value*	Enter a dot1p list number and value. The list number range is from 0 to 7. The range is from 0 to 3.

Defaults

For each dot1p Priority, the default CoS queue value is:

- dot1p Priority: 0 1 2 3 4 5 6 7
- CoS Queue: 0 0 0 1 2 3 3 3

Command Modes

CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 8.3.16.0	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To apply dot1p-queue-mapping, use the `service-class dynamic dot1p` command.

Related Commands

show qos dot1p-queue-mapping — displays the dot1p priority to queue mapping on the switch.

service-class bandwidth-percentage

Specify a minimum bandwidth for queues.

Syntax

```
service-class bandwidth-percentage queue0 number queue1 number
queue2 number queue3 number
```

Parameters	*number*	Enter the bandwidth-weight, as a percentage. The value must be a power of 2. The range is from 1 to 100.
Defaults	none	
Command Modes	CONFIGURATION	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

Guarantee a minimum bandwidth to different queues globally using the `service-class bandwidth-weight` command from CONFIGURATION mode. The command is applied in the same way as the `bandwidth-weight` command in an output QoS policy. The `bandwidth-weight` command in QOS-POLICY-OUT mode supersedes the `service-class bandwidth-weight` command.

When you enable ETS, the egress QoS features in the output QoS policy-map (such as `service-class bandwidth-percentage` and `bandwidth-percentage`), the default bandwidth allocation ratio for egress queues are superseded by ETS configurations. This is to provide compatibility with DCBX. Therefore, Dell Networking recommends disabling ETS when you wish to apply these features exclusively. After you disable ETS on an interface, the configured parameters are applied.

Policy-Based QoS Commands

Policy-based traffic classification is handled with class maps. These maps classify unicast traffic into one of four classes. The system allows you to match multiple class maps and specify multiple match criteria. Policy-based QoS is not supported on logical interfaces, such as port-channels, VLANs, or Loopbacks.

bandwidth-percentage

Assign a percentage of weight to the class/queue.

Syntax `bandwidth-percentage` *`percentage`*

To remove the bandwidth percentage, use the `no bandwidth-percentage` command.

Parameters	*percentage*	Enter the percentage assignment of weight to the class/ queue. The range is from 0 to 100% (granularity 1%).
Defaults	none	
Command Modes	CONFIGURATION (conf-qos-policy-out)	

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	The unit of bandwidth percentage is 1%. A bandwidth percentage of 0 is allowed and disables the scheduling of that class. If the sum of the bandwidth percentages given to all eight classes exceeds 100%, the bandwidth percentage automatically scales down to 100%.	
Related Commands	qos-policy-output — creates a QoS output policy.	

clear qos statistics

Clears Matched Packets, Matched Bytes, and Dropped Packets.

Syntax `clear qos statistics` *interface-name*

Parameters

interface-name	Enter one of the following keywords: • For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.

Defaults none

Command Modes

- EXEC
- EXEC Privilege

Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	When you issue this command, statistical information stored regarding QoS clears and resets to 0. You can access these statistics using the `show qos statistics` command in EXEC mode. When the traffic pattern matches the QoS classification criteria flows, the corresponding counters increment.	
Related Commands	show qos statistics — displays the QoS statistics.	

description

Add a description to the selected policy map or QoS policy.

Syntax

`description {description}`

To remove the description, use the `no description {description}` command.

Parameters

description	Enter a description to identify the policies (80 characters maximum).

Defaults

none

Command Modes

CONFIGURATION (policy-map-input and policy-map-output; conf-qos-policy-in and conf-qos-policy-out; wred)

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Related Commands

policy-map-input — creates an input policy map.

policy-map-output — creates an output policy map.

qos-policy-input — creates an input QoS-policy on the router.

qos-policy-output — creates an output QoS-policy on the router.

wred-profile — creates a WRED profile.

policy-aggregate

Allow an aggregate method of configuring per-port QoS via policy maps. An aggregate QoS policy is part of the policy map (input/output) applied on an interface.

Syntax

`policy-aggregate qos-policy-name`

To remove a policy aggregate configuration, use the `no policy-aggregate qos-policy-name` command.

Parameters

qos-policy-name	Enter the name of the policy map in character format (32 characters maximum).

Defaults

none

Command Modes

CONFIGURATION (policy-map-input and policy-map-output)

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

Aggregate input/output QoS policy applies to all the port ingoing/outgoing traffic. Aggregate input/output QoS policy can coexist with per queue input/output QoS policies.

1. If only aggregate input QoS policy exists, input traffic conditioning configurations (rate-police) apply. Any marking configurations in aggregate input QoS policy are ignored.
2. If aggregate input QoS policy and per class input QoS policy coexist, aggregate input QoS policy preempts per class input QoS policy on input traffic conditioning (rate-police). In other words, if rate police configuration exists in the aggregate QoS policy, the rate police configurations in per class QoS are ignored. Marking configurations in per class input QoS policy still apply to each queue.

Related Commands

policy-map-input — creates an input policy map.

policy-map-output — creates an output policy map.

policy-map-output

Create an output policy map.

Syntax

`policy-map-output` *`policy-map-name`*

To remove a policy map, use the `no policy-map-output` *`policy-map-name`* command.

Parameters

policy-map-name	Enter the name for the policy map in character format (16 characters maximum).

Defaults

none

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To assign traffic to different flows using QoS policy, use the Output Policy map. This command enables Policy-Map-Output Configuration mode (conf-policy-map-out).

Related Commands

service-queue — assigns a class map and QoS policy to different queues.

policy-aggregate — allows an aggregate method of configuring per-port QoS using policy maps.

service-policy output — applies an output policy map to the selected interface.

qos-policy-output

Create a QoS output policy.

Syntax

`qos-policy-output` *`qos-policy-name`*

To remove an existing output QoS policy, use the `no qos-policy-output` *`qos-policy-name`* command.

Parameters

qos-policy-name	Enter your output QoS policy name in character format (32 characters maximum).

Defaults

none

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To specify the name of the output QoS policy, use this command. After the output policy is specified, rate-limit, bandwidth-percentage, and WRED can be defined. This command enables Qos-Policy-Output Configuration mode — (conf-qos-policy-out).

When changing a service-queue configuration in a QoS policy map, all QoS rules are deleted and re-added automatically to ensure that the order of the rules is maintained. As a result, the Matched Packets value shown in the `show qos statistics`command is reset.

Related Commands

bandwidth-percentage — assigns weight to the class/queue percentage.

wred — assigns yellow or green drop precedence.

rate police

Police the incoming traffic rate on the selected interface.

Syntax

`rate police [kbps]` *`committed-rate [burst-KB]`* `[peak [kbps]` *`peak-rate [burst-KB]`*`]`

Parameters

kbps	Enter the keyword `kbps` to specify the rate limit in Kilobits per second (Kbps). Make the following value a multiple of 64. The range is from 0 to 40000000. The default granularity is Megabits per second (Mbps).
committed-rate	Enter the bandwidth in Mbps. The range is from 0 to 10000.

	burst-KB	(OPTIONAL) Enter the burst size in KB. The range is from 16 to 200000. The default is **100**.
	peak *peak-rate*	(OPTIONAL) Enter the keyword `peak` then a number to specify the peak rate in Mbps. The range is from 0 to 10000. The default is the same as designated for *`committed-rate`*.

Defaults Burst size is 100 KB. *`peak-rate`* is the same as *`committed-rate`*. Granularity for *`committed-rate`* and *`peak-rate`* is Mbps unless you use the *`kbps`* option.

Command Modes INTERFACE

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Related Commands rate police — specifies traffic policing on the selected interface.

qos-policy-input — creates a QoS output policy.

rate shape

Shape the traffic output on the selected interface.

Syntax `rate shape [kbps] rate [burst-KB]`

Parameters

kbps	Enter the keyword `kbps` to specify the rate limit in Kilobits per second (Kbps). Make the following value a multiple of 64. The range is from 0 to 40000000. The default granularity is Megabits per second (Mbps).
rate	Enter the outgoing rate in multiples of 10 Mbps. The range is from 10 to 10000.
burst-KB	(OPTIONAL) Enter the burst size in KB. The range is from 0 to 10000. The default is **50**.

Defaults Granularity for rate is **Mbps** unless you use the `kbps` option.

Command Modes QOS-POLICY-OUT

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information When you apply `rate-shape` in QoS policy both on the Queue Level and in Aggregate mode, the queue-based shaping occurs first then aggregate rate shaping.

Related Commands	rate shape — shapes traffic output as part of the designated policy.

service-policy output

Apply an output policy map to the selected interface.

Syntax	`service-policy output` *`policy-map-name`* To remove the output policy map from the interface, use the `no service-policy output` *`policy-map-name`* command.	
Parameters	***policy-map-name***	Enter the name for the policy map in character format (16 characters maximum). You can identify an existing policy map or name one that does not yet exist.
Defaults	none	
Command Modes	INTERFACE	
Command History	**Version 9.2(0.0)**	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.
Usage Information	A single policy-map can be attached to one or more interfaces to specify the service-policy for those interfaces. A policy map attached to an interface can be modified.	
Related Commands	policy-map-output — creates an output policy map.	

service-queue

Assign a class map and QoS policy to different queues.

Syntax	`service-queue` *`queue-id`* `[class-map` *`class-map-name`*`] [qos-policy` *`qos-policy-name`*`]` To remove the queue assignment, use the `no service-queue` *`queue-id`* `[class-map` *`class-map-name`*`] [qos-policy` *`qos-policy-name`*`]` command.	
Parameters	***queue-id***	Enter the value used to identify a queue. The range is from 0 to 3 (four queues per interface; four queues are reserved for control traffic).
	class-map ***class-map-name***	(OPTIONAL) Enter the keyword `class-map` then the class map name assigned to the queue in character format (32 character maximum).

NOTE: This option is available under `policy-map-input` only.

qos-policy ***qos-policy-name***	(OPTIONAL) Enter the keywords `qos-policy` then the QoS policy name assigned to the queue in text format (32 characters maximum). This specifies the input QoS policy assigned to the queue under `policy-map-input` and output QoS policy under `policy-map-output` context.

Defaults none

Command Modes CONFIGURATION (conf-policy-map-in and conf-policy-map-out)

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information There are four queues per interface on the MXL switch. This command assigns a class map or QoS policy to different queues.

Related Commands

class-map — identifies the class map.

service-policy input — applies an input policy map to the selected interface.

service-policy output — applies an output policy map to the selected interface.

set

Mark outgoing traffic with a differentiated service code point (DSCP) or dot1p value.

Syntax `set {ip-dscp value | mac-dot1p value}`

Parameters

ip-dscp ***value***	(OPTIONAL) Enter the keywords `ip-dscp` then the IP DSCP value. The range is from 0 to 63.
mac-dot1p ***value***	Enter the keywords `mac-dot1p` then the dot1p value. The range is from 0 to 7. The allowed values are: 0, 2, 4, 6.

Defaults none

Command Modes CONFIGURATION (conf-qos-policy-in)

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information After the IP DSCP bit is set, other QoS services can then operate on the bit settings.

show qos policy-map

View the QoS policy map information.

Syntax

```
show qos policy-map {summary [interface] | detail [interface]}
```

Parameters

summary *interface*	To view a policy map interface summary, enter the keyword `summary` and optionally one of the following keywords and slot/port or number information: • For a 40 Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information. • For a 10 Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
detail *interface*	To view a policy map interface in detail, enter the keyword `detail` and optionally one of the following keywords and slot/port or number information: • For a 40 Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information. • For a 10 Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.

Defaults none

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example (IPv4)

```
Dell#show qos policy-map detail gigabitethernet 0/0

Interface GigabitEthernet 4/1

Policy-map-input policy
Trust diffserv
Queue# Class-map-name Qos-policy-name
 0       -             q0
 1       CM1           q1
 2       CM2           q2
 3       CM3           q3
 4       CM4           q4
 5       CM5           q5
 6       CM6           q6
 7       CM7           q7
Dell#
```

Example (Summary IPv4)

```
Dell#sho qos policy-map summary

Interface policy-map-input policy-map-output
Gi 4/1     PM1                  -
Gi 4/2     PM2        PMOut
Dell#
```

show qos policy-map-output

View the output QoS policy map details.

Syntax

```
show qos policy-map-output [policy-map-name] [qos-policy-output qos-policy-name]
```

Parameters

policy-map-name	Enter the policy map name.
qos-policy-output *qos-policy-name*	Enter the keyword `qos-policy-output` then the QoS policy name.

Defaults none

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell#show qos policy-map-output

Policy-map-output PolicyMapOutput
Aggregate Qos-policy-name AggPolicyOut
Queue#   Qos-policy-name
  0      qosPolicyOutput
Dell#
```

show qos qos-policy-output

View the output QoS policy details.

Syntax

```
show qos qos-policy-output [qos-policy-name]
```

Parameters

qos-policy-name	Enter the QoS policy name.

Defaults none

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell#show qos qos-policy-output

Dell#show qos qos-policy-output

Qos-policy-output qmap_out
Bandwidth-percentage 10
Qos-policy-output qmap_wg
Rate-shape 100 50
Wred yellow wy
Wred green wg
Dell#
```

show qos statistics

View QoS statistics.

Syntax

```
show qos statistics {wred-profile [interface]} | [interface]
```

Parameters

wred-profile *interface*	Enter the keywords `wred-profile` and optionally one of the following keywords and slot/port or number information: • For a 40–Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information. • For a 10–Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
interface	Enter one of the following keywords and slot/port or number information: • For a 40–Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.

Defaults

none

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Example

```
Dell#show qos statistics
Interface Te 0/20
Queue# Matched Pkts
0       0
1       0
2       0
3       0
Dell#
```

Usage Information

The following describes the `show qos statistics` command in the following example.

Field (ED and EE)	Description
Queue #	Queue Number
Matched Pkts	The number of packets that matched the class-map criteria.

NOTE: When you configure `trust`, matched packet counters are not incremented in this field.

Example

```
Dell#show qos statistics wred-profile
Interface Te 0/20
Drop-statistic  Dropped Pkts
Green           0
Yellow          0
Out of Profile  0
Dell#
```

Usage Information

The following describes the `show qos statistics` command in the following example.

Field (EF)	Description
Queue #	Queue Number
Drop-statistic	Drop statistics for green, yellow, and out-of-profile packets.
Dropped Pkts	The total of the number of packets dropped for green, yellow and out-of-profile.

Related Commands

clear qos statistics — clears counters shown in `show qos statistics`.

show qos wred-profile

View the WRED profile details.

Syntax

```
show qos wred-profile wred-profile-name
```

Parameters

wred-profile-name	Enter the WRED profile name to view the profile details.

Defaults

none

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 8.3.16.1	Introduced on THE MXL 10/40GbE Switch IO Module.

Example

```
Dell#show qos wred-profile

Wred-profile-name min-threshold max-threshold
wred_drop         0             0
wred_ge_y         1024          2048
wred_ge_g         2048          4096
wred_teng_y       4096          8192
wred_teng_g       8192          16384
WRED1             2000          7000
```

wred

Designate the WRED profile to yellow or green traffic.

Syntax

`wred [[{yellow | green} profile-name] ecn]`

To remove the WRED drop precedence, use the `no wred {yellow | green} [profile-name]` command.

Parameters

yellow | green	Enter the keyword `yellow` for yellow traffic. A DSCP value of xxx110 and xxx100 maps to yellow. Enter the keyword green for `green` traffic. A DSCP value of xxx010 maps to green.
profile-name	Enter your WRED profile name in character format (16 character maximum). Or use one of the five pre-defined WRED profile names. Pre-defined Profiles: wred_drop, wred-ge_y, wred_ge_g, wred_teng_y, wred_teng_.
ecn	When you configure `wred ecn <cr>` command, instead of droppping the packets exponentially, Explicit Congestion Notification (ECN) marking is made on the packets.

Defaults

none

Command Modes

CONFIGURATION (conf-qos-policy-out)

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

To assign drop precedence to green or yellow traffic, use this command. If there is no honoring enabled on the input, all the traffic defaults to green drop precedence.

Related Commands

wred-profile — creates a WRED profile and name that profile.

trust — defines the dynamic classification to trust DSCP.

wred-profile

Create a WRED profile and name the profile.

Syntax

`wred-profile` *`wred-profile-name`*

To remove an existing WRED profile, use the `no wred-profile` command.

Parameters

wred-profile-name	Enter your WRED profile name in character format (16 character maximum). Or use one of the pre-defined WRED profile names. You can configure up to 26 WRED profiles plus the five pre-defined profiles, for a total of 31 WRED profiles. Pre-defined Profiles: wred_drop, wred-ge_y, wred_ge_g, wred_teng_y, wred_teng_g.

Defaults

The five pre-defined WRED profiles. When you configure a new profile, the minimum and maximum threshold defaults to predefined wred_ge_g values.

Command Modes

CONFIGURATION

Command History

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.16.1	Introduced on the MXL 10/40GbE Switch IO Module.

Usage Information

Use the default pre-defined profiles or configure your own profile. You cannot delete the pre-defined profiles or their default values. This command enables WRED configuration mode — (conf-wred).

31

reload-type

Configure a switch to reload as a DHCP client in BMP mode with all ports configured for Layer 3 traffic or in Normal mode.

Z9000 S4810 S4820TS6000

Syntax

```
reload-type [bmp | normal—reload {[auto—save {enable | disable}] | [config-scr-download {enable | disable}] | [dhcp-timeout minutes]| retry-count number | vendor-class-identifier description]
```

Use the `disable bmp` command to stop the BMP process.

Parameters

bmp	(Default) Enable the BMP reload type. The system acts as a DHCP client and downloads the FTOS image, configuration and boot files from a specified DHCP server.
normal-reload	Enable the normal reload type and disable BMP reload type. The system retrieves the FTOS image and startup-configuration files from the flash after performing a reload.
auto-save	Configure the auto save option to save the downloaded configuration or script file. They are not saved by default. When auto save is configured, downloaded configurations are automatically saved to the startup configuration. Auto saving the downloaded configurations also requires enabling the **config-scr-download** parameter. Downloaded scripts are automatically saved to the autoexec script.
config-scr-download {enable}	(Optional.) Configure whether the configuration file must be downloaded from the DHCP/file servers (enable).
config-download {disable}	(Optional.) Configure if the downloaded file will not be downloaded from the DHCP/file servers.
dhcp-timeout *minutes*	(Optional) Configure the DHCP timeout (in minutes) after which the BMP reload stops. The range is from 0 to 50. If a range of 0 is entered, the timeout is 0 (no limit). The default is **disabled**.

NOTE:

Dell Networking recommends setting the value to 2 or higher.

retry-count *number*	Configure the number of times to retry loading the FTOS image and configuration download. The retry limit is 0–6. If the retry limit is 0, no retry is performed. The default is 0.
vendor-class-identifier *description*	(Optional) Enter a brief description for DHCP Option 60. Maximum is 64 characters long. NOTE: This parameter replaces the deprecated parameter **user-defined-string**.

Defaults

BMP

Switches running BMP 3.0 reload in BMP mode as a DHCP client with all ports configured for Layer 3 traffic.

Command Modes

GLOBAL CONFIGURATION

Command History

Version 9.0.2.0	Introduced on the S6000.
Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.2(0.0)	Introduced support for `vendor-class-identifier` that replaces deprecated parameter `user-defined-string` Also added support for `retry-count`.
Version 9.1(0.0)	Introduced on the Z9000. Updated the command mode from EXEC Privilege to GLOBAL CONFIGURATION. Updated the parameter from `jumpstart` to `bmp`. Added support for the `config-scr-download` and `user-defined-string` commands.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.1.0	Introduced on the S4810.

Usage Information

For an initial setup, the `config-scr-download` parameter of the `reload-type` command is enabled. After the configuration file is successfully downloaded, the `config—scr-download` parameter is automatically disabled. You can enable it again using the `reload-type` command.

Set the Auto Configuration mode (BMP or Normal reload) using the `reload-type` command. Next, enter the `reload` command to reload the switch in the configured mode.

When a switch reloads in BMP mode, all ports, including the management port, are automatically configured as Layer 3 physical ports. The switch runs DHCP client on all interfaces. You can reconfigure the default startup configuration and DHCP timeout values.

If the switch attempts to contact a DHCP server and one is not found, it enters a loop while reloading in BMP mode. To interrupt the reload and boot up in Normal

mode, enter the `stop bmp` command. The startup configuration is then loaded from the local flash on the switch.

To toggle between Normal and BMP Auto Configuration modes, use the `reload-type` command in BMP 3.0. Reload settings for Auto Configuration mode that you configure are stored in memory and retained for future reboots and BMP software upgrades. To reload the switch in the last configured mode: Normal reload or BMP mode, you can enter the `reload` command at any time.

Upgrade any configuration changes that have changed the NVRAM content by performing a reload on the chassis.

While BMP is on, the Dell Networking OS Command Line Reference Guide prompt changes to "Dell-BMP".

Related Commands

- show reload-type — displays the current reload mode (BMP or Normal mode).
- stop jumpstart — Stops the Jumpstart (BMP) process to prevent a loop if the DHCP server is not found

32

Simple Network Management Protocol (SNMP) and Syslog

This chapter contains commands to configure and monitor the simple network management protocol (SNMP) v1/v2/v3 and Syslog. Both features are supported on the Z-Series S4810 S4820T platform.

The chapter contains the following sections:

- SNMP Commands
- Syslog Commands

SNMP Commands

The following SNMP commands are available in the Dell Networking operating software (FTOS).

The simple network management protocol (SNMP) is used to communicate management information between the network management stations and the agents in the network elements. FTOS supports SNMP versions 1, 2c, and 3, supporting both read-only and read-write modes. FTOS sends SNMP traps, which are messages informing an SNMP management system about the network. FTOS supports up to 16 SNMP trap receivers.

Important Points to Remember

- Typically, 5-second timeout and 3-second retry values on an SNMP server are sufficient for both LAN and WAN applications. If you experience a timeout with these values, the recommended best practice on Dell Networking switches (to accommodate their high port density) is to increase the timeout and retry values on your SNMP server to the following:
  - SNMP Timeout — greater than 3 seconds.
  - SNMP Retry count — greater than 2 seconds.
- If you want to query an E-Series switch using SNMP v1/v2/v3 with an IPv6 address, configure the IPv6 address on a non-management port on the switch.
- If you want to send SNMP v1/v2/v3 traps from an E-Series using an IPv6 address, use a non-management port.
- SNMP v3 informs are not currently supported with IPv6 addresses.
- If you are using access control lists (ACLs) in an SNMP v3 configuration, group ACL overrides user ACL if the user is part of that group.
- SNMP operations are not supported on a virtual local area network (VLAN).

snmp-server enable traps

Enable SNMP traps.

Z-Series S4810 S4820T

Syntax `snmp-server enable traps [notification-type] [notification-option]`

To disable traps, use the `no snmp-server enable traps [notification-type] [notification-option]` command.

Parameters

notification-type

Enter the type of notification from the following list:

- `bgp` — Notification of changes in the BGP process.
- `config` — Notification of changes to the startup or running configuration.
- `ecfm` — Notification of changes to ECFM.
- `ecmp` — Enable an ECMP trap to notify of ECMP or link bundle traffic imbalances.
- `envmon` — For Dell Networking device notifications when an environmental threshold is exceeded.
- `isis` — Notification of intermediate service traps.
- `lacp` — Notification of changes.
- `snmp` — Notification of RFC 1157 traps.
- `stp` — Notification of a state change in the spanning tree protocol (RFC 1493).
- `vlt` — Notification of virtual link trunking.
- `vrrp` — Notification of a state change in a VRRP group.
- `xstp` — Notification of a state change in MSTP (802.1s), RSTP (802.1w), and PVST+.

notification-option

For the `envmon` notification-type, enter one of the following optional parameters:

- `cam-utilization`
- `fan`
- `supply`
- `temperature`

For the `snmp` notification-type, enter one of the following optional parameters:

- `authentication`
- `coldstart`
- `linkdown`
- `linkup`

Defaults

Not enabled.

Command Modes

CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.1(0.0)	Added support for copy-config and ecmp traps.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 8.4.1.0	Added support for VRRP traps.
Version 7.6.1.0	Added support for STP and xSTP traps. Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Usage Information

FTOS supports up to 16 SNMP trap receivers.

For the `cam-utilization` notification option, the system generates syslogs and SNMP traps when the L3 host table or route table utilization goes above the threshold.

If you do not configure this command, no traps controlled by this command are sent. If you do not specify a `notification-type` and `notification-option`, all traps are enabled.

Related Commands

snmp-server community — enables SNMP and sets the community string.

snmp-server host

Configure the recipient of an SNMP trap operation.

Z-Series S4810 S4820T

Syntax

`snmp-server host` *`ip-address | ipv6-address`* `[traps | informs] [version 1 | 2c | 3] [auth | no auth | priv]` [*`community-string`*] `[udp-port` *`port-number`*] [*`notification-type`*]

To remove the SNMP host, use the `no snmp-server host` *`ip-address`* `[traps | informs] [version 1 | 2c | 3] [auth | noauth | priv]` [*`community-string`*] `[udp-port` *`number`*] [*`notification-type`*] command.

Parameters

ip-address	Enter the keyword `host` then the IP address of the host (configurable hosts is limited to 16).
ipv6-address	Enter the keyword host then the IPv6 address of the host in the x:x:x:x::x format.

NOTE: The :: notation specifies successive hexadecimal fields of zero.

traps (OPTIONAL) Enter the keyword `traps` to send trap notifications to the specified host. The default is **traps**.

informs (OPTIONAL) Enter the keyword `informs` to send inform notifications to the specified host. The default is **traps**.

version 1 | 2c | 3 (OPTIONAL) Enter the keyword version to specify the security model then the security model version number `1`, `2c`, or `3`:

- Version `1` is the least secure version.
- Version `3` is the most secure of the security modes.
- Version `2c` allows transmission of informs and counter 64, which allows for integers twice the width of what is normally allowed.

The default is version **1**.

auth (OPTIONAL) Enter the keyword `auth` to specify authentication of a packet without encryption.

noauth (OPTIONAL) Enter the keyword `noauth` to specify no authentication of a packet.

priv (OPTIONAL) Enter the keyword `priv` to specify both authentication and then scrambling of the packet.

community-string Enter a text string (up to 20 characters long) as the name of the SNMP community.

NOTE: For version 1 and version 2c security models, this string represents the name of the SNMP community. The string can be set using this command; however, Dell Networking recommends setting the community string using the `snmp-server community` command before executing this command. For version 3 security model, this string is the USM user security name.

udp-port ***port-number*** (OPTIONAL) Enter the keywords `udp-port` followed by the port number of the remote host to use. The range is from 0 to 65535. The default is **162**.

notification-type (OPTIONAL) Enter one of the following keywords for the type of trap to be sent to the host:

- `bgp` — BGP state change.
- `config` — copy-configuration traps.
- `ecmp` — ecmp and link bundling traffic imbalance traps.
- `envmon` — Environment monitor trap.
- `snmp` — SNMP notification (RFC 1157).
- `stp` — Spanning tree protocol notification (RFC 1493).

- `vrrp` — State change in a VRRP group.
- `xstp` — State change in MSTP (802.1s), RSTP (802.1w), and PVST+.

The default is all trap types are sent to host.

Defaults As above.

Command Modes CONFIGURATION

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.1(0.0)	Added support for config and ecmp traps.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 8.4.1.0	Added support for VRRP traps.
Version 7.6.1.0	Added support for STP and xSTP notification types. Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Usage Information In order to configure the router to send SNMP notifications, enter at least `one snmp-server host` command. If you enter the command with no keywords, all trap types are enabled for the host. If you do not enter an `snmp-server host` command, no notifications are sent.

In order to enable multiple hosts, issue a separate `snmp-server host` command for each host. You can specify multiple notification types in the command for each host.

When multiple `snmp-server host` commands are given for the same host and type of notification (trap or inform), each succeeding command overwrites the previous command. Only the last `snmp-server host` command will be in effect. For example, if you enter an `snmp-server host inform` command for a host and then enter another `snmp-server host inform` command for the same host, the second command replaces the first command.

The `snmp-server host` command is used with the `snmp-server enable` command. Use the `snmp-server enable` command to specify which SNMP notifications are sent globally. For a host to receive most notifications, at least one

`snmp-server enable` command and the `snmp-server host` command for that host must be enabled.

NOTE: For v1 / v2c trap configuration, if the community-string is not defined using the `snmp-server community` command prior to using this command, the default form of the `snmp-server community` command automatically is configured with the community-name the same as specified in the `snmp-server host` command.

Configuring Informs

To send an inform, use the following steps:

1. Configure a remote engine ID.
2. Configure a remote user.
3. Configure a group for this user with access rights.
4. Enable traps.
5. Configure a host to receive informs.

Related Commands

snmp-server enable traps — enables SNMP traps.

snmp-server community — configures a new community SNMPv1 or SNMPv2c.

Syslog Commands

The following commands allow you to configure logging functions on all Dell Networking switches.

clear logging

Clear the messages in the logging buffer.

Z-Series S4810 S4820T

Syntax `clear logging`

Defaults none

Command Modes EXEC Privilege

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.

	Version 7.5.1.0	Introduced on the C-Series.
	E-Series legacy command	

Related Commands show logging — displays logging settings and system messages in the internal buffer.

logging

Configure an IP address or host name of a Syslog server where logging messages are sent. Multiple logging servers of both IPv4 and/or IPv6 can be configured.

Z-Series S4810 S4820T

Syntax `logging {ip-address | ipv6-address | hostname}`

To disable logging, use the `no logging` command.

Parameters

ip-address	Enter the IPv4 address in dotted decimal format.
ipv6-address	Enter the IPv6 address in the x:x:x:x::X format.
	NOTE: The :: notation specifies successive hexadecimal fields of zeros.
hostname	Enter the name of a host already configured and recognized by the switch.

Defaults Disabled.

Command Modes CONFIGURATION

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 8.4.1.0	Added support for IPv6.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Usage Information	Multiple logging servers of both IPv4 and/or IPv6 can be configured.
Related Commands	logging on — enables the logging asynchronously to logging buffer, console, Syslog server, and terminal lines. logging trap — enables logging to the Syslog server based on severity.

logging buffered

Enable logging and specify which messages are logged to an internal buffer. By default, all messages are logged to the internal buffer.

Z-Series S4810 S4820T

Syntax

`logging buffered [level] [size]`

To return to the default values, use the `default logging buffered` command.

To disable logging stored to an internal buffer, use the `no logging buffered` command.

Parameters

level	(OPTIONAL) Indicate a value from 0 to 7 or enter one of the following equivalent words: `emergencies`, `alerts`, `critical`, `errors`, `warnings`, `notifications`, `informational`, or `debugging`. The default is **7** or **debugging**.
size	(OPTIONAL) Indicate the size, in bytes, of the logging buffer. The number of messages buffered depends on the size of each message. The range is from 40960 to 524288. The default is **40960 bytes**.

Defaults level = **7**; size = **40960 bytes**

Command Modes CONFIGURATION

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.

	E-Series legacy command
Usage Information	When you decrease the buffer size, all messages stored in the buffer are lost. Increasing the buffer size does not affect messages stored in the buffer.
Related Commands	clear logging — clears the logging buffer. default logging buffered — returns the logging buffered parameters to the default setting. show logging — displays the logging setting and system messages in the internal buffer.

logging console

Specify which messages are logged to the console.

Z-Series S4810S4820T

Syntax

`logging console [level]`

To return to the default values, use the `default logging console` command.

To disable logging to the console, use the `no logging console` command.

Parameters

level	(OPTIONAL) Indicate a value from 0 to 7 or enter one of the following parameters: `emergencies`, `alerts`, `critical`, `errors`, `warnings`, `notifications`, `informational`, or `debugging`. The default is **7** or **debugging**.

Defaults

level = **7**; size = **debugging**

Command Modes

CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Related Commands	clear logging — clears the logging buffer. default logging console — returns the logging console parameters to the default setting. show logging — displays the logging setting and system messages in the internal buffer.

logging monitor

Specify which messages are logged to Telnet applications.

Z-Series S4810 S4820T

Syntax	`logging monitor [level]` To disable logging to terminal connections, use the `no logging monitor` command.
Parameters	***level*** — Indicate a value from 0 to 7 or enter one of the following parameters: `emergencies`, `alerts`, `critical`, `errors`, `warnings`, `notifications`, `informational`, or `debugging`. The default is **7** or **debugging**.
Defaults	**7** or **debugging**
Command Modes	CONFIGURATION
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*. The following is a list of the FTOS version history for this command. **Version 9.2(0.0)** — Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only. **Version 8.3.19.0** — Introduced on the S4820T. **Version 8.3.11.1** — Introduced on the Z9000. **Version 8.3.7.0** — Introduced on the S4810. **Version 7.6.1.0** — Introduced on the S-Series. **Version 7.5.1.0** — Introduced on the C-Series. **E-Series legacy command**
Related Commands	default logging monitor — returns the logging monitor parameters to the default setting.

logging source-interface

Specify that the IP address of an interface is the source IP address of Syslog packets sent to the Syslog server.

Z9000 S4810 S4820T

Syntax

`logging source-interface` *interface*

To disable this command and return to the default setting, use the `no logging source-interface` command.

Parameters

interface — Enter the following keywords and slot/port or number information:

- For Loopback interfaces, enter the keyword `loopback` then a number from zero (0) to 16383.
- For the management interface on the RPM, enter the keyword `ManagementEthernet` then the slot/port information. The slot range is from 0 to 1 and the port range is 0.
- For a Port Channel interface, enter the keywords `port-channel` then a number. Tthe range is from 1 to 128.
- For a ten-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For VLAN interface, enter the keyword `vlan` then a number from 1 to 4094.

Defaults

Not configured.

Command Modes

CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 8.5.1.0	Added support for 4-port 40G line cards on ExaScale.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.

	E-Series legacy command
Usage Information	Syslog messages contain the IP address of the interface used to egress the router. By configuring the `logging source-interface` command, the Syslog packets contain the IP address of the interface configured.
Related Commands	logging — enables logging to the Syslog server.

show logging

Display the logging settings and system messages logged to the internal buffer of the switch.

Z-Series S4810 S4820T

Syntax

```
show logging [number | history [reverse][number] | reverse [number] | summary]
```

Parameters

number	(OPTIONAL) Enter the number of messages displayed in the output. The range is from 1 to 65535.
history	(OPTIONAL) Enter the keyword `history` to view only information in the Syslog history table.
reverse	(OPTIONAL) Enter the keyword `reverse` to view the Syslog messages in FIFO (first in, first out) order.
summary	(OPTIONAL) Enter the keyword `summary` to view a table showing the number of messages per type and per slot. Slots *7* and *8* represent RPMs.

Command Modes

- EXEC
- EXEC Privilege

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.

E-Series legacy command

Example (Partial)

```
FTOS#show logging
Syslog logging: enabled
    Console logging: level debugging
    Monitor logging: level debugging
    Buffer logging: level debugging, 5604 Messages Logged,
Size (524288 bytes)
    Trap logging: level informational
Oct 8 09:25:37: %RPM1:RP1 %BGP-5-ADJCHANGE: Connection with
neighbor 223.80.255.254 closed. Hold time
expired
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor
192.200.13.2 Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor
192.1.1.13 Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 1.1.14.2
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor
192.1.1.14 Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 1.1.11.2
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 192.1.1.5
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 192.4.1.3
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 192.1.1.4
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 192.1.1.6
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor
192.1.1.12 Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor
192.1.1.15 Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 192.1.1.3
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor
192.200.12.2 Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor 1.1.10.2
Up
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Session closed by
neighbor 1.1.10.2 (Hold time expired)
Oct 8 09:25:38: %RPM1:RP1 %BGP-5-ADJCHANGE: Neighbor
192.200.14.7 Up
Oct 8 09:26:25: %RPM1:RP1 %BGP-5-ADJCHANGE: Connection with
neighbor 1.1.11.2 closed. Neighbor recycled
Oct 8 09:26:25: %RPM1:RP1 %BGP-5-ADJCHANGE: Connection with
neighbor 1.1.14.2 closed. Neighbor recycled
--More--
```

Example (History)

```
FTOS#show logging history
Syslog History Table: 1 maximum table entries,
saving level Warnings or higher
  SNMP notifications not Enabled
%RPM:0:0 %CHMGR-2-LINECARDDOWN - Line card 3 down - IPC timeout
FTOS#
```

show logging driverlog stack-unit

Display the driver log for the specified stack member.

S4810 S4820T Z9000

Syntax	`show logging driverlog stack-unit` *unit#*
Parameters	**stack-unit** *unit#* — Enter the keywords `stack-unit` followed by the stack member ID of the switch for which you want to display the driver log. The range is from 0 to 1.
defaults	none
Command Modes	• EXEC • EXEC Privilege
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*. The following is a list of the FTOS version history for this command. **Version 9.2(0.0)** — Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only. **Version 8.3.19.0** — Introduced on the S4820T. **Version 8.3.11.1** — Introduced on the Z9000. **Version 8.3.7.0** — Introduced on the S4810. **Version 7.6.1.0** — Introduced on the S-Series.
Usage Information	This command displays internal software driver information, which may be useful during troubleshooting switch initialization errors, such as a downed Port-Pipe.

terminal monitor

Configure the FTOS to display messages on the monitor/terminal.

Z-Series S4810 S4820T

Syntax	`terminal monitor` To return to default settings, use the `terminal no monitor` command.
defaults	Disabled.
Command Modes	• EXEC • EXEC Privilege

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
E-Series legacy command	

Related Commands

logging monitor — sets the logging parameters on the monitor/terminal.

33

Storm Control

The Dell Networking operating software (FTOS) storm control feature allows you to limit or suppress traffic during a traffic storm (Broadcast/Unknown Unicast Rate Limiting or Multicast on the C-Series and S-Series).

Storm control is supported on the Dell Networking Z-Series S4810 S4820T platforms.

Important Points to Remember

- Interface commands can only be applied on physical interfaces (virtual local area networks [VLANs] and link aggregation group [LAG] interfaces are not supported).
- An INTERFACE-level command only supports storm control configuration on ingress.
- An INTERFACE-level command overrides any CONFIGURATION-level ingress command for that physical interface, if both are configured.
- You can apply the CONFIGURATION-level storm control commands at ingress or egress and are supported on all physical interfaces.
- When storm control is applied on an interface, the percentage of storm control applied is calculated based on the advertised rate of the line card. It is not based on the speed setting for the line card.
- Do not apply per-VLAN quality of service (QoS) on an interface that has storm control enabled (either on an interface or globally).
- When you enable broadcast storm control on an interface or globally on ingress, and DSCP marking for a DSCP value 1 is configured for the data traffic, the traffic goes to queue 1 instead of queue 0.
- Similarly, if you enable unicast storm control on an interface or globally on ingress, and DSCP marking for a DSCP value 2 is configured for the data traffic, the traffic goes to queue 2 instead of queue 0.

NOTE: Bi-directional traffic (unknown unicast and broadcast) along with egress storm control causes the configured traffic rates split between the involved ports. The percentage of traffic that each port receives after the split is not predictable. These ports can be in the same/different port pipes or the same/different line cards.

show storm-control unknown-unicast

Display the storm control unknown-unicast configuration.

Z-Series S4810 S4820TS6000

Syntax `show storm-control unknown-unicast [interface]`

Parameters

interface	(OPTIONAL) Enter one of the following interfaces to display the interface specific storm control configuration: • For a 1-Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then y the slot/port information. • For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.

- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/ port information.

Defaults none

Command Modes
- EXEC
- EXEC Privilege

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.0.2.0	Introduced on the S6000.
Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.5.1.0	Added support for 4-port 40G line cards on ExaScale.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
Version 6.5.1.0	Introduced on the E-Series.

storm-control broadcast (Configuration)

Configure the percentage of broadcast traffic allowed in the network.

Z-Series S4810 S4820TS6000

Syntax `storm-control broadcast [packets_per_second in]`

To disable broadcast rate-limiting, use the `no storm-control broadcast [packets_per_second in]` command.

storm-control broadcast `[percentage decimal_value` **in | out]**`| [wred-profile name]]`**[**`packets_per_second` **in]**

To disable broadcast rate-limiting, use the **storm-control broadcast** [*percentage decimal_value* in | out] | [wred-profile name]] [*packets_per_second* **in]** command.

Parameters

percentagedecimal_value **in | out**	Enter the percentage of broadcast traffic allowed in or out of the network. Optionally, you can designate a decimal value percentage, for example, 55.5%. The decimal range is from .1 to .9.

	wred-profile name	Enter the keyword **wred-profile** followed by the profile name to designate a wred-profile.
	***packets_per_second* in**	Enter the packets per second of broadcast traffic allowed into the network. The range is from 0 to 33554368.

Defaults none

Command Modes CONFIGURATION (conf)

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
Version 7.4.1.0	E-Series Only: Added the `percentage decimal value` option.
Version 6.5.1.0	Introduced on the E-Series.

Usage Information Broadcast storm control is valid on Layer 2/Layer 3 interfaces only. Layer 2 broadcast traffic is treated as unknown-unicast traffic.

storm-control multicast (Configuration)

Configure the packets per second (pps) of multicast traffic allowed into the C-Series and S-Series networks only.

Z-SeriesS4810 S4820TS6000

Syntax `storm-control multicast` *`packets_per_second`* `in`

To disable storm-control for multicast traffic into the network, use the `no storm-control multicast` *`packets_per_second`* `in` command.

Parameters

***packets_per_second* in**	Enter the packets per second of multicast traffic allowed into the network. The range is from 0 to 33554368.

Defaults none

Command Modes	CONFIGURATION (conf)
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*. The following is a list of the Dell Networking OS version history for this command.

Version 9.0.2.0	Introduced on the S6000.
Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the C-series and S-Series.

Usage Information	Broadcast traffic (all 0xFs) should be counted against the broadcast storm control meter, not against the multicast storm control meter. It is possible, however, that some multicast control traffic may get dropped when storm control thresholds are exceeded.

storm-control broadcast (Interface)

Configure the percentage of broadcast traffic allowed on an interface (ingress only).

Z-Series S4810 S4820TS6000

Syntax	`storm-control broadcast [packets_per_second in]` To disable broadcast storm control on the interface, use the `no storm-control broadcast [packets_per_second in]` command.
Parameters	***packets_per_second in*** Enter the packets per second of broadcast traffic allowed into the network. The range is from 0 to 33554368.
Defaults	none
Command Modes	INTERFACE (conf-if-*interface-slot/port*)
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*. The following is a list of the Dell Networking OS version history for this command.

Version 9.0.2.0	Introduced on the S6000.
Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.11.1	Introduced on the Z9000.
Version 8.3.7.0	Introduced on the S4810.
Version 7.6.1.0	Introduced on the S-Series.
Version 7.5.1.0	Introduced on the C-Series.
Version 7.4.1.0	E-Series Only: Added the `percentage decimal value` option.
Version 6.5.1.0	Introduced on the E-Series.

34

Uplink Failure Detection (UFD)

Uplink failure detection (UFD) provides detection of the loss of upstream connectivity and, if you use this with NIC teaming, automatic recovery from a failed link. UFD is supported on the S4810 S4820T Dell Networking platform.

clear ufd-disable

Re-enable one or more downstream interfaces on the switch/router that are in a UFD-Disabled Error state so that an interface can send and receive traffic.

S4810 S4820T

Syntax `clear ufd-disable {interface` *`interface`* `| uplink-state-group` *`group-id`*`}`

Parameters

interface *interface* — Specify one or more downstream interfaces. For *interface*, enter one of the following interface types:

- Fast Ethernet: `fastethernet {`*`slot/port | slot/port-range`*`}`
- 1 Gigabit Ethernet: `gigabitethernet {`*`slot/port | slot/ port-range`*`}`
- 10 Gigabit Ethernet: `tengigabitethernet {`*`slot/ port |slot/ port-range`*`}`
- Port channel: `port-channel {`*`1-512 | port-channel-range`*`}`

Where `port-range` and `port-channel-range` specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example: `gigabitethernet 1/1-2,5,9,11-12 port-channel 1-3,5`. A comma is required to separate each port and port-range entry.

uplink-state-group *group-id* — Re-enables all UFD-disabled downstream interfaces in the group. The valid group-id values are from 1 to 16.

Defaults A downstream interface in a UFD-disabled uplink-state group is also disabled and is in a UFD-Disabled Error state.

Command Modes CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Related Commands

- downstream — assigns a port or port-channel to the uplink-state group as a downstream interface.
- uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

debug uplink-state-group

Enable debug messages for events related to a specified uplink-state group or all groups.

S4810 S4820T

Syntax

`debug uplink-state-group [group-id]`

To turn off debugging event messages, enter the `no debug uplink-state-group [group-id]` command.

Parameters

group-id	Enables debugging on the specified uplink-state group. The valid group-id values are from 1 to 16.

Defaults

none

Command Modes

EXEC Privilege

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.

	Version 8.3.12.0	Introduced on the S4810.
	Version 8.4.2.3	Introduced on the S-Series S50.

Related Commands clear ufd-disable — re-enables downstream interfaces that are in a UFD-Disabled Error state.

description

Enter a text description of an uplink-state group.

S4810 S4820T

Syntax `description text`

Parameters

text	Text description of the uplink-state group. The maximum length is 80 alphanumeric characters.

Defaults none

Command Modes UPLINK-STATE-GROUP

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Example

```
FTOS(conf-uplink-state-group-16)# description test
FTOS(conf-uplink-state-group-16)#
```

Related Commands uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

downstream

Assign a port or port-channel to the uplink-state group as a downstream interface.

S4810 S4820T

Syntax

`downstream` *interface*

To delete an uplink-state group, enter the `no downstream interface` command.

Parameters

interface — Enter one of the following interface types:

- Fast Ethernet: `fastethernet {`*slot/port* | *slot/port-range*`}`
- 1 Gigabit Ethernet: `gigabitethernet {`*slot/port* | *slot/port-range*`}`
- 10 Gigabit Ethernet: `tengigabitethernet {`*slot/port* |*slot/port-range*`}`
- Port channel: `port-channel {`*1-512* | *port-channel-range*`}`

Where `port-range` and `port-channel-range` specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example: `gigabitethernet 1/1-2,5,9,11-12 port-channel 1-3,5`. A comma is required to separate each port and port-range entry.

Defaults

none

Command Modes

UPLINK-STATE-GROUP

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Usage Information

You can assign physical port or port-channel interfaces to an uplink-state group.

You can assign an interface to only one uplink-state group. Configure each interface assigned to an uplink-state group as either an upstream or downstream interface, but not both.

You can assign individual member ports of a port channel to the group. An uplink-state group can contain either the member ports of a port channel or the port channel itself, but not both.

Related Commands

- upstream — assigns a port or port-channel to the uplink-state group as an upstream interface.
- uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

downstream auto-recover

Enable auto-recovery so that UFD-disabled downstream ports in an uplink-state group automatically come up when a disabled upstream port in the group comes back up.

S4810 S4820T

Syntax

`downstream auto-recover`

To disable auto-recovery on downstream links, use the `no downstream auto-recover` command.

Defaults

The auto-recovery of UFD-disabled downstream ports is enabled.

Command Modes

UPLINK-STATE-GROUP

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Related Commands

- downstream — assigns a port or port-channel to the uplink-state group as a downstream interface.
- uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

downstream disable links

Configure the number of downstream links in the uplink-state group that are disabled if one upstream link in an uplink-state group goes down.

S4810 S4820T

Syntax

`downstream disable links {number |all}`

To revert to the default setting, use the `no downstream disable links` command.

Parameters

number	Enter the number of downstream links to be brought down by UFD. The range is from 1 to 1024.
all	Brings down all downstream links in the group.

Defaults

No downstream links are disabled when an upstream link in an uplink-state group goes down.

Command Modes

UPLINK-STATE-GROUP

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Usage Information

A user-configurable number of downstream interfaces in an uplink-state group are put into a link-down state with an UFD-Disabled error message when one upstream interface in an uplink-state group goes down.

If all upstream interfaces in an uplink-state group go down, all downstream interfaces in the same uplink-state group are put into a link-down state.

Related Commands

- downstream — assigns a port or port-channel to the uplink-state group as a downstream interface.
- uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

enable

Enable uplink state group tracking for a specific UFD group.

S4810 S4820T

Syntax

`enable`

To disable upstream-link tracking without deleting the uplink-state group, use the `no enable` command.

Defaults

Upstream-link tracking is automatically enabled in an uplink-state group.

Command Modes

UPLINK-STATE-GROUP

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Related Commands

- uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

show running-config uplink-state-group

Display the current configuration of one or more uplink-state groups.

S4810 S4820T

Syntax

`show running-config uplink-state-group [group-id]`

Parameters

group-id	Displays the current configuration of all uplink-state groups or a specified group. The valid group-id values are from 1 to 16.

Defaults

none

Command Modes

- EXEC

- EXEC Privilege

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Example

```
FTOS#show running-config uplink-state-group
!
no enable
uplink state track 1
downstream GigabitEthernet 0/2,4,6,11-19
upstream TengigabitEthernet 0/48, 52
upstream PortChannel 1
!
uplink state track 2
downstream GigabitEthernet 0/1,3,5,7-10
upstream TengigabitEthernet 0/56,60
```

Related Commands

- show uplink-state-group — displays the status information on a specified uplink-state group or all groups.
- uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

show uplink-state-group

Display status information on a specified uplink-state group or all groups.

S4810 S4820T

Syntax

```
show uplink-state-group [group-id] [detail]
```

Parameters

group-id	Displays status information on a specified uplink-state group or all groups. The valid group-id values are from 1 to 16.
detail	Displays additional status information on the upstream and downstream interfaces in each group

Defaults

none

Command Modes

- EXEC
- EXEC Privilege

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Example

```
FTOS# show uplink-state-group
Uplink State Group: 1 Status: Enabled, Up
Uplink State Group: 3 Status: Enabled, Up
Uplink State Group: 5 Status: Enabled, Down
Uplink State Group: 6 Status: Enabled, Up
Uplink State Group: 7 Status: Enabled, Up
Uplink State Group: 16 Status: Disabled, Up

FTOS# show uplink-state-group 16
Uplink State Group: 16 Status: Disabled, Up

FTOS#show uplink-state-group detail
(Up): Interface up (Dwn): Interface down (Dis): Interface
disabled
Uplink State Group    : 1 Status: Enabled, Up
Upstream Interfaces   :
Downstream Interfaces :

Uplink State Group    : 3 Status: Enabled, Up
Upstream Interfaces   : Gi 0/46(Up) Gi 0/47(Up)
Downstream Interfaces : Te 13/0(Up) Te 13/1(Up) Te 13/3(Up) Te
13/5(Up) Te 13/6(Up)

Uplink State Group    : 5 Status: Enabled, Down
Upstream Interfaces   : Gi 0/0(Dwn) Gi 0/3(Dwn) Gi 0/5(Dwn)
Downstream Interfaces : Te 13/2(Dis) Te 13/4(Dis) Te
13/11(Dis) Te 13/12(Dis) Te 13/13(Dis) Te 13/14(Dis) Te
13/15(Dis)

Uplink State Group    : 6 Status: Enabled, Up
Upstream Interfaces   :
Downstream Interfaces :

Uplink State Group    : 7 Status: Enabled, Up
Upstream Interfaces   :
Downstream Interfaces :

Uplink State Group    : 16 Status: Disabled, Up
Upstream Interfaces   : Gi 0/41(Dwn) Po 8(Dwn)
Downstream Interfaces : Gi 0/40(Dwn)
```

Related Commands

- show running-config uplink-state-group — displays the current configuration of one or more uplink-state groups.
- uplink-state-group — create an uplink-state group and enables the tracking of upstream links.

uplink-state-group

Create an uplink-state group and enable the tracking of upstream links on a switch/ router.

S4810 S4820T

Syntax

`uplink-state-group` *`group-id`*

To delete an uplink-state group, enter the `no uplink-state-group group-id` command.

Parameters

group-id	Enter the ID number of an uplink-state group. The range is from 1 to 16.

Defaults

none

Command Modes

CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Usage Information

After you enter the command, to assign upstream and downstream interfaces to the group, enter Uplink-State-Group Configuration mode.

An uplink-state group is considered to be operationally up if at least one upstream interface in the group is in the Link-Up state.

An uplink-state group is considered to be operationally down if no upstream interfaces in the group are in the Link-Up state. No uplink-state tracking is performed when a group is disabled or in an operationally down state.

To disable upstream-link tracking without deleting the uplink-state group, use the `no enable` command in uplink-state-group configuration mode.

Example

```
FTOS(conf)#uplink-state-group 16
FTOS(conf)#
02:23:17: %RPM0-P:CP %IFMGR-5-ASTATE_UP: Changed uplink state
group Admin state to up: Group 16
```

Related Commands

- show running-config uplink-state-group — displays the current configuration of one or more uplink-state groups.
- show uplink-state-group — displays the status information on a specified uplink-state group or all groups.

upstream

Assign a port or port-channel to the uplink-state group as an upstream interface.

S4810 S4820T

Syntax

`upstream` *interface*

To delete an uplink-state group, use the `no upstream interface` command.

Parameters

interface — Enter one of the following interface types:

- Fast Ethernet: `fastethernet {`*slot/port | slot/port-range*`}`
- 1 Gigabit Ethernet: `gigabitethernet {`*slot/port | slot/port-range*`}`
- 10 Gigabit Ethernet: `tengigabitethernet {`*slot/port | slot/port-range*`}`
- 40 Gigabit Ethernet: `fortyGigE {`*slot/port | slot/port-range*`}`
- Port channel: `port-channel {`*1-512 | port-channel-range*`}`

Where `port-range` and `port-channel-range` specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example: `gigabitethernet 1/1-2,5,9,11-12 port-channel 1-3,5`. A comma is required to separate each port and port-range entry.

Defaults

none

Command Modes

UPLINK-STATE-GROUP

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.
Version 8.4.2.3	Introduced on the S-Series S50.

Usage Information

You can assign physical port or port-channel interfaces to an uplink-state group.

You can assign an interface to only one uplink-state group. Configure each interface assigned to an uplink-state group as either an upstream or downstream interface, but not both.

You can assign individual member ports of a port channel to the group. An uplink-state group can contain either the member ports of a port channel or the port channel itself, but not both.

Example

```
FTOS(conf-uplink-state-group-16)# upstream gigabitethernet
1/10-15
FTOS(conf-uplink-state-group-16)#
```

Related Commands

- downstream — assigns a port or port-channel to the uplink-state group as a downstream interface.
- uplink-state-group — creates an uplink-state group and enables the tracking of upstream links.

35

Virtual Link Trunking (VLT)

Virtual link trunking (VLT) is supported on the Z-Series S4810S4820T platform.

VLT allows physical links between two chassis to appear as a single virtual link to the network core. VLT eliminates the requirement for Spanning Tree protocols by allowing link aggregation group (LAG) terminations on two separate distribution or core switches, and by supporting a loop-free topology. VLT provides Layer 2 multipathing, creating redundancy through increased bandwidth and enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

NOTE: When you launch the VLT link, the VLT peer-ship is not established if any of the following is **TRUE**:

- The VLT System-MAC configured on both the VLT peers do not match.
- The VLT Unit-Id configured on both the VLT peers are identical.
- The VLT System-MAC or Unit-Id is configured only on one of the VLT peers.
- The VLT domain ID is not the same on both peers.

If the VLT peer-ship is already established, changing the System-MAC or Unit-Id does not cause VLT peer-ship to go down.

Also, if the VLT peer-ship is already established and the VLT Unit-Id or System-MAC are configured on both peers, then changing the CLI configurations on the VLT Unit-Id or System-MAC is rejected if any of the following become **TRUE**:

- After making the CLI configuration change, the VLT Unit-Id becomes identical on both peers.
- After making the CLI configuration change, the VLT System-MAC do not match on both peers.

When the VLT peer-ship is already established, you can remove the VLT Unit-Id or System-MAC configuration from either or both peers. However, removing configuration settings can cause the VLT ports to go down if you configure the Unit-Id or System-MAC on only one of the VLT peers.

back-up destination

Configure the IPv4 or IPv6 address of the management interface on the remote VLT peer to be used as the endpoint of the VLT backup link for sending out-of-band hello messages.

Z9000 S4810 S4820T

Syntax

```
back-up destination {[ipv4-address] | [ipv6 ipv6-address]
[interval seconds]}
```

Parameters

ipv4–address	Enter the IPv4 address of the backup destination.
ipv6	Enter the keyword `ipv6` then an IPv6 address in the X:X:X:X::X format.

	interval *seconds*	Enter the keyword `interval` to specify the time interval to send hello messages. The range is from 1 to 5 seconds. The default is 1 second.

Defaults **1 second**

Command Modes VLT DOMAIN

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.2(0.2)	Added support for IPv6.
Version 9.0.0.0	Introduced on the Z9000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.8.0	Introduced on the S4810.

clear vlt statistics

Clear the statistics on VLT operations.

Z9000 S4810 S4820T

Syntax
```
clear vlt statistics [arp | domain | igmp-snoop | mac |
multicast | ndp]
```

Parameters

domain	Clear the VLT statistics for the domain.
multicast	Display the VLT statistics for multicast.
mac	Clear the VLT statistics for the MAC address.
arp	Clear the VLT statistics for ARP.
igmp-snoop	Clear the VLT statistics for IGMP snooping.
ndp	Clear the VLT statistics for NDP.

Command Modes EXEC

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.2(0.2)	Added `multicast` and `ndp` parameters.
Version 9.0.0.0	Introduced on the Z9000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Introduced on the S4810.

Example

```
VLT ARP Statistics
----------------
ARP Tunnel Pkts sent:0
ARP Tunnel Pkts Rcvd:0
ARP-sync Pkts Sent:0
ARP-sync Pkts Rcvd:0
ARP Reg Request sent:19
ARP Reg Request rcvd:10
```

Related Commands show vlt statistics — displays statistics on VLT operations.

delay-restore

Configure the delay in bringing up VLT ports after reload or peer-link restoration between the VLT peer switches.

Z-Series S4810 S4820T

Syntax `delay-restore`

Parameters

delay-restore	Enter the amount of time, in seconds, to delay bringing up the VLT ports after the VLTi device is reloaded or after the peer-link is restored between VLT peer switches. The range from 1 to 1200. The default is **90 seconds**.

Defaults Not configured.

Command Modes VLT DOMAIN

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.0.0	Introduced on the Z9000.

Version 8.3.19.0	Introduced on the S8420T.
Version 8.3.12.0	Introduced on the S4810.

Usage Information

To delay the system from bringing up the VLT port for a brief period to allow IGMP Snooping and Layer 3 routing protocols to converge, use the `delay-restore` parameter. Use this feature:

- after a VLT device is reloaded.
- if the Peer VLT device was up at the time the VLTi link failed to the time when it was restored.

Related Commands

show vlt statistics — displays statistics on VLT operations.

lacp ungroup member-independent

Prevent possible loop during the bootup of a VLT peer switch or a device that accesses the VLT domain.

Z-Series S4810 S4820T

Syntax

```
lacp ungroup member-independent {vlt | port-channel}
```

Parameters

port-channel	Force all LACP port-channel members to become switchports.
vlt	Force all VLT LACP members to become switchports.

Defaults

Not configured.

Command Modes

CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.0.0	Introduced on the Z9000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Added port-channel parameter on the S4810.
Version 8.3.8.0	Introduced on the S4810.

Usage Information

LACP on the VLT ports (on a VLT switch or access device), which are members of the virtual link trunk, is not brought up until the VLT domain is recognized on the access device.

On the S4810, during boot-up in a stacking configuration, the system must be able to reach the DHCP server with the boot image and configuration image. During boot-up, only untagged DHCP requests are sent to the DHCP server to receive an offer on static LAGs between switches. The DHCP server must be configured to start in BMP mode. If switches are connected using LACP port-channels like the VLT peer and Top of Rack (ToR), use the **port-channel** parameter on the ToR-side configuration to allow member ports of an ungrouped LACP port-channel to inherit vlan membership of that port channel to ensure untagged packets that are sent by a VLT peer device reach the DHCP server located on the ToR.

To ungroup the VLT and port-channel configurations, use the **no lacp ungroup member independent** command on a VLT port channel, depending on whether the port channel is VLT or non-VLT.

Example

```
FTOS (conf)#lacp ungroup member-independent ?
port-channel                 LACP port-channel members become
switchports
vlt                                    All VLT LACP members
become switchports
```

peer-link port-channel

Configure the specified port channel as the chassis interconnect trunk between VLT peers in the domain.

Z-Series S4810 S4820T

Syntax

`peer-link port-channel` *`port-channel-number`* `{peer-down-vlan` *`vlan id`*`}`

Parameters

port-channel-number	Enter the port-channel number that acts as the interconnect trunk.
peer-down-vlan *vlan id*	Enter the keyword `peer-down-vlan` then a VLAN ID to configure the VLAN that the VLT peer link uses when the VLT peer is down.

Defaults

Not configured.

Command Modes

VLT DOMAIN

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

Version 9.0.0.0	Introduced on the Z9000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.12.0	Added support for the **peer-down-vlan** parameter.
Version 8.3.8.0	Introduced on the S4810.

Usage Information

To configure the VLAN from where the VLT peer forwards packets received over the VLTi from an adjacent VLT peer that is down, use the **peer-down-vlan** parameter. When a VLT peer with bare metal provisioning (BMP) is booting up, it sends untagged DHCP discover packets to its peer over the VLTi. To ensure that the DHCP discover packets are forwarded to the VLAN that has the DHCP server, use this configuration.

primary-priority

Assign the priority for master election among VLT peers.

S4810 S4820T

Syntax

```
[no] primary-priority
```

Parameters

value	To configure the primary role on a VLT peer, enter a lower value than the priority value of the remote peer. The range is from 1 to 65535.

Default 32768

Command Modes VLT DOMAIN

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.8.0	Introduced on the S4810.

Usage Information

After you configure the VLT domain on each peer switch on both sides of the interconnect trunk, by default, the FTOS software elects a primary and secondary VLT peer device. To reconfigure the primary role of VLT peer switches, use the `priority` command.

show vlt mismatch

Display mismatches in VLT parameters.

Z9000 S4810 S4820TS6000

Syntax `show vlt mismatch`

Command Modes EXEC

Command History This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.

The following is a list of the Dell Networking OS version history for this command.

Version 9.0.2.0	Introduced on the S6000.
Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.2(0.2)	Introduced on the Z9000, S4810, and S4820T.

Example

```
Dell#show vlt mismatch
Domain
-----------
Parameters        Local              Peer
------------     -----------       ---------
    Unit-ID       0                   15

Vlan-config
------------
Vlan-ID      Local Mode     Peer Mode
---------   ------------ -----------

     100              --                  L3

Vlan IPV4 Multicast Status
--------------------------
  Vlan-ID       Local Status       Peer Status
  --------      ------------       -----------

  4094               Active              Inactive

Dell#
```

system-mac

Reconfigure the default MAC address for the domain.

Z-Series S4810 S4820T

Syntax `system-mac` *mac-address*

Parameters	**mac-address**	Enter the system MAC address for the VLT domain.
Defaults	Not configured.	
Command Modes	VLT DOMAIN	
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *Dell Networking OS Command Line Reference Guide*.	
	The following is a list of the Dell Networking OS version history for this command.	
	Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 9.0.0.0	Introduced on the Z9000.
	Version 8.3.19.0	Introduced on the S4820T.
	Version 8.3.8.0	Introduced on the S4810.
Usage Information	When you create a VLT domain on a switch, Dell Networking OS automatically creates a VLT-system MAC address used for internal system operations.	
	To reconfigure the default MAC address for the domain by entering a new MAC address in the format nn:nn:nn:nn:nn:nn, use the `system-mac` command.	
	You must also reconfigure the same MAC address on the VLT peer switch.	

unit-id

Explicitly configure the default unit ID of a VLT peer switch.

Z-Series S4810S4820T

Syntax	`unit-id [0 | 1]`	
Parameters	**0 | 1**	Configure the default unit ID of a VLT peer switch. Enter `0` for the first peer or enter `1` for the second peer.
Defaults	Automatically assigned based on the MAC address of each VLT peer. The peer with the lower MAC address is assigned unit 0; the peer with the higher MAC address is assigned unit 1.	
Command Modes	VLT DOMAIN	
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.	
	The following is a list of the FTOS version history for this command.	
	Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.

Version 9.0.0.0	Introduced on the Z9000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.8.0	Introduced on the S4810.

Usage Information

When you create a VLT domain on a switch, FTOS automatically assigns a unique unit ID (0 or 1) to each peer switch. The unit IDs are used for internal system operations. Use the `unit-id` command to explicitly configure the unit ID of a VLT peer. Configure a different unit ID (0 or 1) on each peer switch.

To minimize the time required for the VLT system to determine the unit ID assigned to each peer switch when one peer reboots, use this command.

vlt domain

Enable VLT on a switch, configure a VLT domain, and enter VLT-domain configuration mode.

Z-Series S4810 S4820T

Syntax

```
vlt domain domain-id
```

Parameters

domain-id	Enter the Domain ID number. Configure the same domain ID on the peer switch. The range of domain IDs is from 1 to 1000.

Command Modes

CONFIGURATION

Command History

This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.

The following is a list of the FTOS version history for this command.

Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
Version 9.0.0.0	Introduced on the Z9000.
Version 8.3.19.0	Introduced on the S4820T.
Version 8.3.8.0	Introduced on the S4810.

Usage Information

The VLT domain ID must be the same between the two VLT devices. If the domain ID is not the same, a syslog message is generated and VLT does not launch.

Related Commands

show vlt — uses the `show vlt brief` command to display the delay-restore value.

vlt-peer-lag port-channel

Associate the port channel to the corresponding port channel in the VLT peer for the VLT connection to an attached device.

Z-Series S4810 S4820T

Syntax	`vlt-peer-lag port-channel` *id-number*	
Parameters	***id-number***	Enter the respective vlt port-channel number of the peer device.
Defaults	Not configured.	
Command Modes	INTERFACE PORT-CHANNEL	
Command History	This guide is platform-specific. For command information about other platforms, refer to the relevant *FTOS Command Line Reference Guide*.	
	The following is a list of the FTOS version history for this command.	
	Version 9.2(0.0)	Introduced on the M I/O Aggregator. This command is supported in Programmable-Mux (PMUX) mode only.
	Version 9.0.0.0	Introduced on the Z9000.
	Version 8.3.19.0	Introduced on the S4820T.
	Version 8.3.8.0	Introduced on the S4810.

Overview

VLT allows physical links between two chassis to appear as a single virtual link to the network core or other switches such as Edge, Access, or top-of-rack (ToR).

VLT reduces the role of spanning tree protocols (STPs) by allowing link aggregation group (LAG) terminations on two separate distribution or core switches, and by supporting a loop-free topology. (To prevent the initial loop that may occur prior to VLT being established, use a spanning tree protocol. After VLT is established, you may use rapid spanning tree protocol (RSTP) to prevent loops from forming with new links that are incorrectly connected and outside the VLT domain.)

VLT provides Layer 2 multipathing, creating redundancy through increased bandwidth, enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

Virtual link trunking offers the following benefits:

- Allows a single device to use a LAG across two upstream devices.
- Eliminates STP-blocked ports.
- Provides a loop-free topology.
- Uses all available uplink bandwidth.

- Provides fast convergence if either the link or a device fails.
- Optimized forwarding with virtual router redundancy protocol (VRRP).
- Provides link-level resiliency.
- Assures high availability.

CAUTION: Dell Networking does not recommend enabling Stacking and VLT simultaneously. If you enable both features at the same time, unexpected behavior occurs.

As shown in the following example, VLT presents a single logical Layer 2 domain from the perspective of attached devices that have a virtual link trunk terminating on separate chassis in the VLT domain. However, the two VLT chassis are independent Layer2/Layer3 (L2/L3) switches for devices in the upstream network. L2/L3 control plane protocols and system management features function normally in VLT mode. Features such as VRRP and internet group management protocol (IGMP) snooping require state information coordinating between the two VLT chassis. IGMP and VLT configurations must be identical on both sides of the trunk to ensure the same behavior on both sides.

The following example shows VLT deployed on S4810 S4820T switches. The S4810 S4820T switches appear as a single virtual switch from the point of view of the switch or server supporting link aggregation control protocol (LACP).

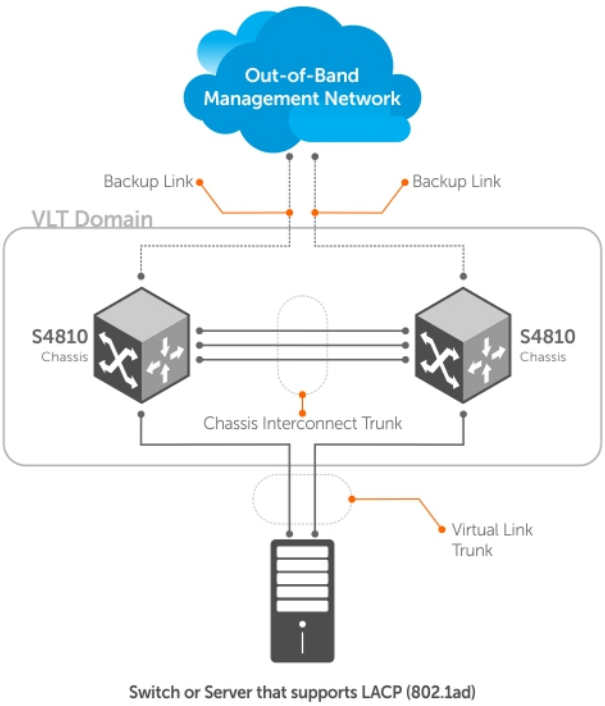

Figure 3. VLT on S4810 S4820T Switches

VLT on Core Switches

You can also deploy VLT on core switches.

Uplinks from servers to the access layer and from access layer to the aggregation layer are bundled in LAG groups with end-to-end Layer 2 multipathing. This set up requires “horizontal” stacking at the access layer and VLT at the aggregation layer such that all the uplinks from servers to access and access to

aggregation are in Active-Active Load Sharing mode. This example provides the highest form of resiliency, scaling, and load balancing in data center switching networks.

The following example shows stacking at the access, VLT in aggregation, and Layer 3 at the core.

The aggregation layer is mostly in the L2/L3 switching/routing layer. For better resiliency in the aggregation, Dell Networking recommends running the internal gateway protocol (IGP) on the VLTi VLAN to synchronize the L3 routing table across the two nodes on a VLT system.

Enhanced VLT

An enhanced VLT (eVLT) configuration creates a port channel between two VLT domains by allowing two different VLT domains, using different VLT domain ID numbers, connected by a standard link aggregation control protocol (LACP) LAG to form a loop-free Layer 2 topology in the aggregation layer.

This configuration supports a maximum of four units, increasing the number of available ports and allowing for dual redundancy of the VLT. The following example shows how the core/aggregation port density in the Layer 2 topology is increased using eVLT. For inter-VLAN routing and other Layer 3 routing, you need a separate Layer 3 router.

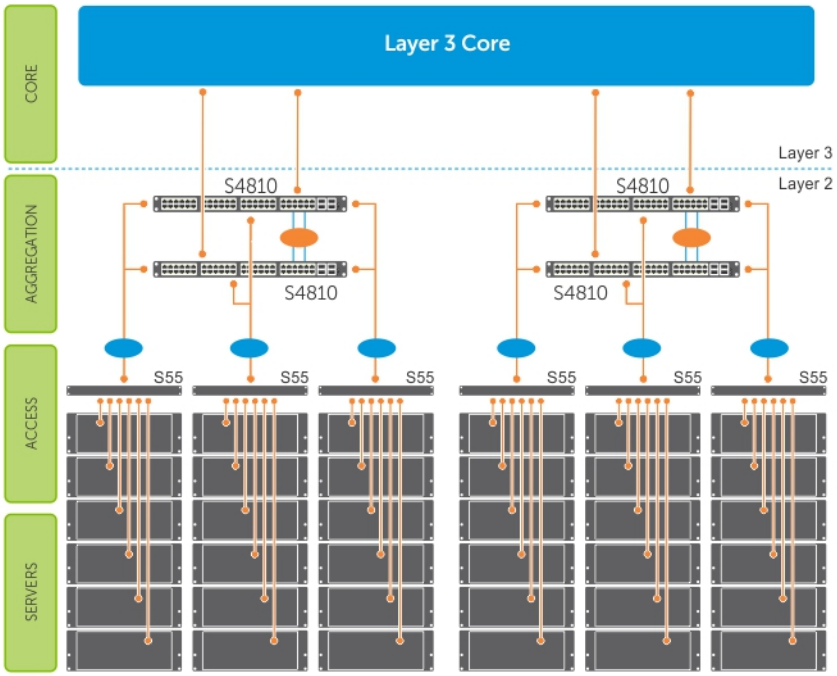

Figure 4. Enhanced VLT

VLT Terminology

The following are key VLT terms.

- **Virtual link trunk (VLT)** — The combined port channel between an attached device and the VLT peer switches.
- **VLT backup link** — The backup link monitors the vitality of VLT peer switches. The backup link sends configurable, periodic keep alive messages between the VLT peer switches.
- **VLT interconnect (VLTi)** — The link used to synchronize states between the VLT peer switches. Both ends must be on 10G or 40G interfaces.

- **VLT domain** — This domain includes both the VLT peer devices, VLT interconnect, and all of the port channels in the VLT connected to the attached devices. It is also associated to the configuration mode that you must use to assign VLT global parameters.
- **VLT peer device** — One of a pair of devices that are connected with the special port channel known as the VLT interconnect (VLTi).

VLT peer switches have independent management planes. A VLT interconnect between the VLT chassis maintains synchronization of L2/L3 control planes across the two VLT peer switches. The VLT interconnect uses either 10G or 40G user ports on the chassis.

A separate backup link maintains heartbeat messages across an out-of-band (OOB) management network. The backup link ensures that node failure conditions are correctly detected and are not confused with failures of the VLT interconnect. VLT ensures that local traffic on a chassis does not traverse the VLTi and takes the shortest path to the destination via directly attached links.

Configure Virtual Link Trunking

VLT requires that you enable the feature and then configure the same VLT domain, backup link, and VLT interconnect on both peer switches.

Important Points to Remember

- You cannot enable S4810 S4820T stacking simultaneously with VLT. If you enable both at the same time, unexpected behavior occurs. Refer to VLT and Stacking.
- VLT port channel interfaces must be switch ports.
- If you include RSTP on the system, configure it before VLT. Refer to Configure Rapid Spanning Tree.
- Dell Networking strongly recommends that the VLTi (VLT interconnect) be a static LAG and that you disable LACP on the VLTi.
- Ensure that the spanning tree root bridge is at the Aggregation layer. If you enable RSTP on the VLT device, refer to RSTP and VLT for guidelines to avoid traffic loss.
- If you reboot both VLT peers in BMP mode and the VLT LAGs are static, the DHCP server reply to the DHCP discover offer may not be forwarded by the ToR to the correct node. To avoid this scenario, configure the VLT LAGs to the ToR and the ToR port channel to the VLT peers with LACP. If supported by the ToR, enable the `lacp-ungroup` feature on the ToR using the `lacp ungroup member-independent port-channel` command.
- If the `lacp-ungroup` feature is not supported on the ToR, reboot the VLT peers one at a time. After rebooting, verify that VLTi (ICL) is active before attempting DHCP connectivity.
- When you enable IGMP snooping on the VLT peers, ensure the value of the `delay-restore` command is not less than the query interval.
- When you enable Layer 3 routing protocols on VLT peers, make sure the delay-restore timer is set to a value that allows sufficient time for all routes to establish adjacency and exchange all the L3 routes between the VLT peers before you enable the VLT ports.
- Only use the `lacp ungroup member-independent` command if the system connects to nodes using bare metal provisioning (BMP) to upgrade or boot from the network.
- Ensure that you configure all port channels where LACP ungroup is applicable as hybrid ports and as untagged members of a VLAN. BMP uses untagged dynamic host configuration protocol (DHCP) packets to communicate with the DHCP server.
- If the DHCP server is located on the ToR and the VLTi (ICL) is down due to a failed link when a VLT node is rebooted in BMP mode, it is not able to reach the DHCP server, resulting in BMP failure.
- If the source is connected to an orphan (non-spanned, non-VLT) port in a VLT peer, the receiver is connected to a VLT (spanned) port-channel, and the VLT port-channel link between the VLT peer connected to the source and TOR is down, traffic is duplicated due to route inconsistency between peers. To avoid this scenario, Dell Networking recommends configuring both the source and the receiver on a spanned VLT VLAN.

- Bulk Sync happens only for Global IPv6 Neighbors; Link-local neighbor entries are not synced.
- If all of the following conditions are true, MAC addresses may not be synced correctly:
  - VLT peers use VLT interconnect (VLTi)
  - Sticky MAC is enabled on an orphan port in the primary or secondary peer
  - MACs are currently inactive

  If this scenario occurs, use the `clear mac-address-table sticky all` command on the primary or secondary peer to correctly sync the MAC addresses.
- If static ARP is enabled on only one VLT peer, entries may be overwritten during bulk sync.
- For multiple VLT LAGs configured on the same VLAN, if a host is learned on one VLT LAG and there is a station move between LAGs, the link local address redirects to the VLTi link on one of the peers. If this occurs, clear the link local address that is redirecting to the VLTi link.

Configuration Notes

When you configure VLT, the following conditions apply.

- VLT domain
  - A VLT domain supports two chassis members, which appear as a single logical device to network access devices connected to VLT ports through a port channel.
  - A VLT domain consists of the two core chassis, the interconnect trunk, backup link, and the LAG members connected to attached devices.
  - Each VLT domain has a unique MAC address that you create or VLT creates automatically.
  - ARP tables are synchronized between the VLT peer nodes.
  - VLT peer switches operate as separate chassis with independent control and data planes for devices attached on non-VLT ports.
  - One chassis in the VLT domain is assigned a primary role; the other chassis takes the secondary role. The primary and secondary roles are required for scenarios when connectivity between the chassis is lost. VLT assigns the primary chassis role according to the lowest MAC address. You can configure the primary role.
  - In a VLT domain, the peer switches must run the same Dell Networking operating system (FTOS) software version.
  - Separately configure each VLT peer switch with the same VLT domain ID and the VLT version. If the system detects mismatches between VLT peer switches in the VLT domain ID or VLT version, the VLT Interconnect (VLTi) does not activate. To find the reason for the VLTi being down, use the `show vlt statistics` command to verify that there are mismatch errors, then use the `show vlt brief` command on each VLT peer to view the VLT version on the peer switch. If the VLT version is more than one release different from the current version in use, the VLTi does not activate.
  - The chassis members in a VLT domain support connection to orphan hosts and switches that are not connected to both switches in the VLT core.
- VLT interconnect (VLTi)
  - The VLT interconnect must consist of either 10G or 40G ports. A maximum of eight 10G or four 40G ports is supported. A combination of 10G and 40G ports is not supported.
  - A VLT interconnect over 1G ports is *not* supported.
  - The port channel must be in Default mode (not Switchport mode) to have VLTi recognize it.
  - The system automatically includes the required VLANs in VLTi. You do not need to manually select VLANs.
  - VLT peer switches operate as separate chassis with independent control and data planes for devices attached to non-VLT ports.

- Port-channel link aggregation (LAG) across the ports in the VLT interconnect is required; individual ports are not supported. Dell Networking strongly recommends configuring a static LAG for VLTi.
- The VLT interconnect synchronizes L2 and L3 control-plane information across the two chassis.
- The VLT interconnect is used for data traffic only when there is a link failure that requires using VLTi in order for data packets to reach their final destination.
- Unknown, multicast, and broadcast traffic can be flooded across the VLT interconnect.
- MAC addresses for VLANs configured across VLT peer chassis are synchronized over the VLT interconnect on an egress port such as a VLT LAG. MAC addresses are the same on both VLT peer nodes.
- ARP entries configured across the VLTi are the same on both VLT peer nodes.
- If you shut down the port channel used in the VLT interconnect on a peer switch in a VLT domain in which you did not configure a backup link, the switch's role displays in the `show vlt brief` command output as Primary instead of Standalone.
- When you change the default VLAN ID on a VLT peer switch, the VLT interconnect may flap.
- In a VLT domain, the following software features are supported on VLTi: link layer discovery protocol (LLDP), flow control, port monitoring, jumbo frames, and data center bridging (DCB).
- When you enable the VLTi link, the link between the VLT peer switches is established if the following configured information is true on both peer switches:

  * the VLT system MAC address matches.
  * the VLT unit-id is not identical.

  NOTE: If you configure the VLT system MAC address or VLT unit-id on only one of the VLT peer switches, the link between the VLT peer switches is not established. Each VLT peer switch must be correctly configured to establish the link between the peers.

- If the link between the VLT peer switches is established, changing the VLT system MAC address or the VLT unit-id causes the link between the VLT peer switches to become disabled. However, removing the VLT system MAC address or the VLT unit-id may disable the VLT ports if you happen to configure the unit ID or system MAC address on only one VLT peer at any time.
- If the link between VLT peer switches is established, any change to the VLT system MAC address or unit-id fails if the changes made create a mismatch by causing the VLT unit-ID to be the same on both peers and/or the VLT system MAC address does not match on both peers.
- If you replace a VLT peer node, preconfigure the switch with the VLT system MAC address, unit-id, and other VLT parameters before connecting it to the existing VLT peer switch using the VLTi connection.

- VLT backup link

  - In the backup link between peer switches, heartbeat messages are exchanged between the two chassis for health checks. The default time interval between heartbeat messages over the backup link is 1 second. You can configure this interval. The range is from 1 to 5 seconds. DSCP marking on heartbeat messages is CS6.
  - In order that the chassis backup link does not share the same physical path as the interconnect trunk, Dell Networking recommends using the management ports on the chassis and traverse an out-of-band management network. The backup link can use user ports, but not the same ports the interconnect trunk uses.
  - The chassis backup link does not carry control plane information or data traffic. Its use is restricted to health checks only.

- Virtual link trunks (VLTs) between access devices and VLT peer switches

  - To connect servers and access switches with VLT peer switches, you use a VLT port channel, as shown in Overview. Up to 48 port-channels are supported; up to eight member links are supported in each port channel between the VLT domain and an access device.
  - The discovery protocol running between VLT peers automatically generates the ID number of the port channel that connects an access device and a VLT switch. The discovery protocol uses LACP

properties to identify connectivity to a common client device and automatically generates a VLT number for port channels on VLT peers that connects to the device. The discovery protocol requires that an attached device always runs LACP over the port-channel interface.

- VLT provides a loop-free topology for port channels with endpoints on different chassis in the VLT domain.
- VLT uses shortest path routing so that traffic destined to hosts via directly attached links on a chassis does not traverse the chassis-interconnect link.
- VLT allows multiple active parallel paths from access switches to VLT chassis.
- VLT supports port-channel links with LACP between access switches and VLT peer switches. Dell Networking recommends using static port channels on VLTi.
- If VLTi connectivity with a peer is lost but the VLT backup connectivity indicates that the peer is still alive, the VLT ports on the Secondary peer are orphaned and are shut down.

  * In one possible topology, a switch uses the BMP feature to receive its IP address, configuration files, and boot image from a DHCP server that connects to the switch through the VLT domain. In the port-channel used by the switch to connect to the VLT domain, configure the port interfaces on each VLT peer as hybrid ports before adding them to the port channel (refer to Connecting a VLT Domain to an Attached Access Device (Switch or Server)). To configure a port in Hybrid mode so that it can carry untagged, single-tagged, and double-tagged traffic, use the `portmode hybrid` command in Interface Configuration mode as described in Configuring Native VLANs.
  * For example, if the DHCP server is on the ToR and VLTi (ICL) is down (due to either an unavailable peer or a link failure), whether you configured the VLT LAG as static or LACP, when a single VLT peer is rebooted in BMP mode, it cannot reach the DHCP server, resulting in BMP failure.

- Software features supported on VLT port-channels

  - In a VLT domain, the following software features are supported on VLT port-channels: 802.1p, ingress and egress ACLs, BGP, DHCP relay, IS-IS, OSPF, active-active PIM-SM, PIM-SSM, VRRP, Layer 3 VLANs, LLDP, flow control, port monitoring, jumbo frames, IGMP snooping, sFlow, ingress and egress ACLs, and Layer 2 control protocols RSTP only).

    NOTE: PVST+ passthrough is supported in a VLT domain. PVST+ BPDUs does not result in an interface shutdown. PVST+ BPDUs for a nondefault VLAN is flooded out as any other L2 multicast packet. On a default VLAN, RTSP is part of the PVST+ topology in that specific VLAN (default VLAN).
  - For detailed information about how to use VRRP in a VLT domain, refer to the following *VLT and VRRP interoperability* section.
  - For information about configuring IGMP Snooping in a VLT domain, refer to VLT and IGMP Snooping.
  - All system management protocols are supported on VLT ports, including SNMP, RMON, AAA, ACL, DNS, FTP, SSH, Syslog, NTP, RADIUS, SCP, TACACS+, Telnet, and LLDP.
  - Enable Layer 3 VLAN connectivity VLT peers by configuring a VLAN network interface for the same VLAN on both switches.
  - Dell Networking does not recommend enabling peer-routing if the CAM is full. To enable peer-routing, a minimum of two local DA spaces for wild card functionality are required.

- Software features supported on VLT physical ports

  - In a VLT domain, the following software features are supported on VLT physical ports: 802.1p, LLDP, flow control, port monitoring, and jumbo frames.
- Software features not supported with VLT

  - In a VLT domain, the following software features are supported on non-VLT ports: 802.1x, , DHCP snooping, FRRP, IPv6 dynamic routing, ingress and egress QOS.

- VLT and VRRP interoperability
  - In a VLT domain, VRRP interoperates with virtual link trunks that carry traffic to and from access devices (refer to Overview). The VLT peers belong to the same VRRP group and are assigned master and backup roles. Each peer actively forwards L3 traffic, reducing the traffic flow over the VLT interconnect.
  - VRRP elects the router with the highest priority as the master in the VRRP group. To ensure VRRP operation in a VLT domain, configure VRRP group priority on each VLT peer so that a peer is either the master or backup for all VRRP groups configured on its interfaces. For more information, refer to Setting VRRP Group (Virtual Router) Priority.
  - To verify that a VLT peer is consistently configured for either the master or backup role in all VRRP groups, use the `show vrrp` command on each peer.
  - Also configure the same L3 routing (static and dynamic) on each peer so that the L3 reachability and routing tables are identical on both VLT peers. Both the VRRP master and backup peers must be able to locally forward L3 traffic in the same way.
  - In a VLT domain, although both VLT peers actively participate in L3 forwarding as the VRRP master or backup router, the `show vrrp` command output displays one peer as master and the other peer as backup.
- Failure scenarios
  - On a link failover, when a VLT port channel fails, the traffic destined for that VLT port channel is redirected to the VLTi to avoid flooding.
  - When a VLT switch determines that a VLT port channel has failed (and that no other local port channels are available), the peer with the failed port channel notifies the remote peer that it no longer has an active port channel for a link. The remote peer then enables data forwarding across the interconnect trunk for packets that would otherwise have been forwarded over the failed port channel. This mechanism ensures reachability and provides loop management. If the VLT interconnect fails, the VLT software on the primary switch checks the status of the remote peer using the backup link. If the remote peer is up, the secondary switch disables all VLT ports on its device to prevent loops.
  - If all ports in the VLT interconnect fail, or if the messaging infrastructure fails to communicate across the interconnect trunk, the VLT management system uses the backup link interface to determine whether the failure is a link-level failure or whether the remote peer has failed entirely. If the remote peer is still alive (heartbeat messages are still being received), the VLT secondary switch disables its VLT port channels. If keepalive messages from the peer are not being received, the peer continues to forward traffic, assuming that it is the last device available in the network. In either case, after recovery of the peer link or reestablishment of message forwarding across the interconnect trunk, the two VLT peers resynchronize any MAC addresses learned while communication was interrupted and the VLT system continues normal data forwarding.
  - If the primary chassis fails, the secondary chassis takes on the operational role of the primary.
- The SNMP MIB reports VLT statistics.

Primary and Secondary VLT Peers

Primary and Secondary VLT Peers are supported on the Z9000 S4810 S4820T platform.

To prevent issues when connectivity between peers is lost, you can designate Primary and Secondary roles for VLT peers . You can elect or configure the Primary Peer. By default, the peer with the lowest MAC address is selected as the Primary Peer. You can configure another peer as the Primary Peer using the VLT `domain` *`domain-id`* `role priority` *`priority-value`* command.

If the VLTi link fails, the status of the remote VLT Primary Peer is checked using the backup link. If the remote VLT Primary Peer is available, the Secondary Peer disables all VLT ports to prevent loops.

If all ports in the VLTi link fail or if the communication between VLTi links fails, VLT checks the backup link to determine the cause of the failure. If the failed peer can still transmit heartbeat messages, the Secondary Peer disables all VLT member ports and any Layer 3 interfaces attached to the VLAN associated with the VLT domain. If heartbeat messages are not received, the Secondary Peer forwards

traffic assumes the role of the Primary Peer. If the original Primary Peer is restored, the VLT peer reassigned as the Primary Peer retains this role and the other peer must be reassigned as a Secondary Peer. Peer role changes are reported as SNMP traps.

VLT Bandwidth Monitoring

When bandwidth usage of the VLTi (ICL) exceeds 80%, a syslog error message (shown in the following message) and an SNMP trap are generated.

```
%STKUNIT0-M:CP %VLTMGR-6-VLT-LAG-ICL: Overall Bandwidth utilization of VLT-ICL-
LAG (port-channel 25)
crosses threshold. Bandwidth usage (80 )
```

When the bandwidth usage drops below the 80% threshold, the system generates another syslog message (shown in the following message) and an SNMP trap.

```
%STKUNIT0-M:CP %VLTMGR-6-VLT-LAG-ICL: Overall Bandwidth utilization of VLT-ICL-
LAG (port-channel 25)
reaches below threshold. Bandwidth usage (74 )VLT show remote port channel
status
```

VLT and Stacking

You cannot enable stacking on S4810 S4820T units with VLT.

If you enable stacking on a unit on which you want to enable VLT, you must first remove the unit from the existing stack. For information about how to remove a unit from a stack, refer to Removing a Unit from an S-Series Stack. After you remove the unit, you can configure VLT on the unit.

VLT and IGMP Snooping

When configuring IGMP Snooping with VLT, ensure the configurations on both sides of the VLT trunk are identical to get the same behavior on both sides of the trunk.

When you configure IGMP snooping on a VLT node, the dynamically learned groups and multicast router ports are automatically learned on the VLT peer node.

VLT IPv6

VLT IPv6 is supported on the Z9000 S4810 S4820T platform.

The following features have been enhanced to support IPv6:

- **VLT Sync** — Entries learned on the VLT interface are synced on both VLT peers.
- **Non-VLT Sync** — Entries learned on non-VLT interfaces are synced on both VLT peers.
- **Tunneling** — Control information is associated with tunnel traffic so that the appropriate VLT peer can mirror the ingress port as the VLT interface rather than pointing to the VLT peer's VLTi link.
- **Statistics and Counters** — Statistical and counter information displays IPv6 information when applicable.
- **Heartbeat** — You can configure an IPv4 or IPv6 address as a backup link destination. You cannot use an IPv4 and an IPv6 address simultaneously.

VLT Port Delayed Restoration

With FTOS version 8.3.12.0 8.3.12.0 8.3.19.0, when a VLT node boots up, if the VLT ports have been previously saved in the start-up configuration, they are not immediately enabled.

To ensure MAC and ARP entries from the VLT per node are downloaded to the newly enabled VLT node, the system allows time for the VLT ports on the new node to be enabled and begin receiving traffic.

The `delay-restore` feature waits for all saved configurations to be applied, then starts a configurable timer. After the timer expires, the VLT ports are enabled one-by-one in a controlled manner. The delay between bringing up each VLT port-channel is proportional to the number of physical members in the port-channel. The default is 90 seconds.

To change the duration of the configurable timer, use the `delay-restore` command.

If you enable IGMP snooping, IGMP queries are also sent out on the VLT ports at this time allowing any receivers to respond to the queries and update the multicast table on the new node.

This delay in bringing up the VLT ports also applies when the VLTi link recovers from a failure that caused the VLT ports on the secondary VLT peer node to be disabled.

PIM-Sparse Mode Support on VLT

The designated router functionality of the PIM Sparse-Mode multicast protocol is supported on VLT peer switches for multicast sources and receivers that are connected to VLT ports.

VLT peer switches can act as a last-hop router for IGMP receivers and as a first-hop router for multicast sources.

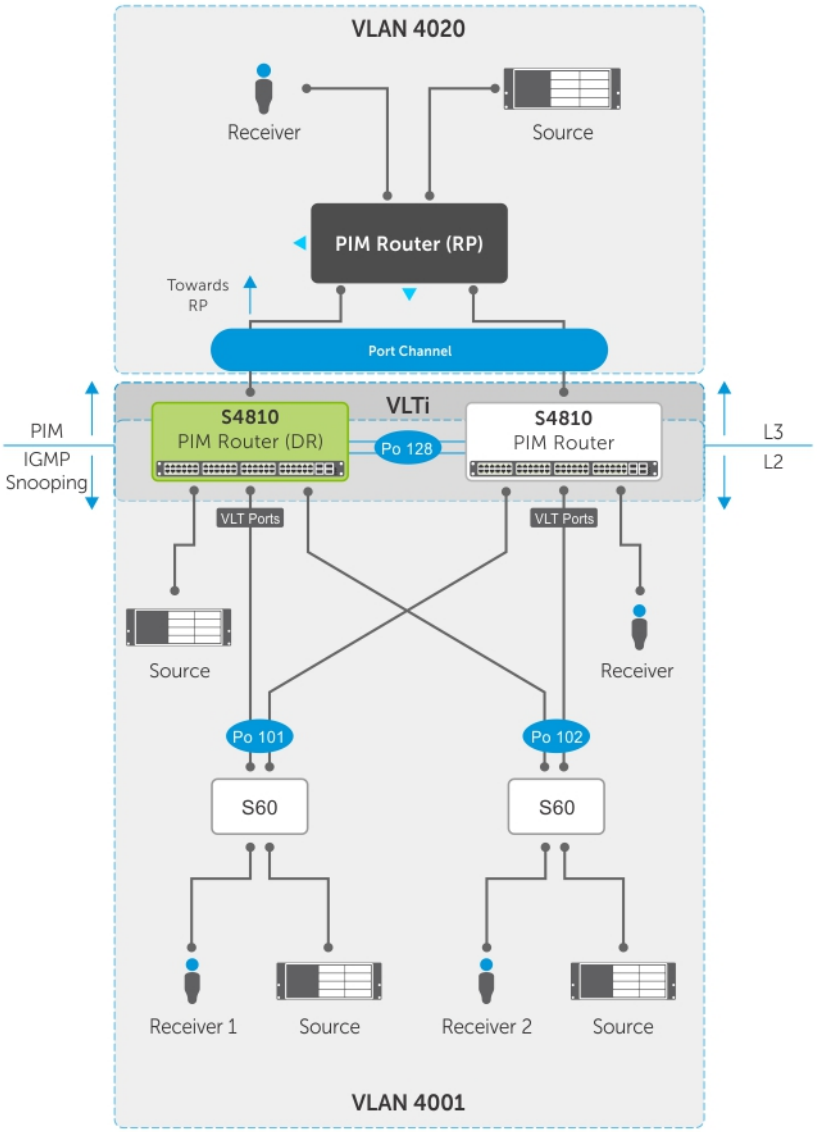

Figure 5. PIM-Sparse Mode Support on VLT

On each VLAN where the VLT peer nodes act as the first hop or last hop routers, one of the VLT peer nodes is elected as the PIM designated router. If you configured IGMP snooping along with PIM on the VLT VLANs, you must configure VLTi as the static multicast router port on both VLT peer switches. This ensures that for first hop routers, the packets from the source are redirected to the designated router (DR) if they are incorrectly hashed. In addition to being first-hop or last -hop routers, the peer node can also act as an intermediate router.

On a VLT-enabled PIM router, if any PIM neighbor is reachable through a Spanned Layer 3 (L3) VLAN interface, this must be the **only** PIM-enabled interface to reach that neighbor. A Spanned L3 VLAN is any L3 VLAN configured on both peers in a VLT domain. This does not apply to server-side L2 VLT ports because they do not connect to any PIM routers. These VLT ports can be members of multiple PIM-enabled L3 VLANs for compatibility with IGMP.

To route traffic to and from the multicast source and receiver, enable PIM on the L3 side connected to the PIM router using the `ip pim sparse-mode` command.

Each VLT peer runs its own PIM protocol independently of other VLT peers. To ensure the PIM protocol states or multicast routing information base (MRIB) on the VLT peers are synced, if the incoming interface (IIF) and outgoing interface (OIF) are Spanned, the multicast route table is synced between the VLT peers.

To verify the PIM neighbors on the VLT VLAN and on the multicast port, use the `show ip pim neighbor`, `show ip igmp snooping mrouter`, and `show running config` commands.

You cannot configure VLT peer nodes as rendezvous points, but you can connect PIM routers to VLT ports.

If the VLT node elected as the designated router fails and you enable VLT Multicast Routing, multicast routes are synced to the other peer for traffic forwarding to ensure minimal traffic loss. If you did not enable VLT Multicast Routing, traffic loss occurs until the other VLT peer is selected as the DR.

VLT Routing

VLT routing is supported on the Z9000 S4810 S4820T platform.

Layer 2 protocols from the ToR to the server are intra-rack and inter-rack. No spanning tree is required, but interoperability with spanning trees at the aggregation layer is supported. Communication is active-active, with no blocked links. MAC tables are synchronized between VLT nodes for bridging and you can enable IGMP snooping.

Because VLT ports are Layer 2 ports and not IP interfaces, VLT Unicast and VLT Multicast routing protocols do not operate directly on VLT ports. You must add the VLT ports as a member of one or more VLANs and assign IP addresses to these VLANs. VLT Unicast and VLT Multicast routing protocols require VLAN IP interfaces for operation. Protocols such as BGP, ISIS, OSPF, and PIM are compatible with VLT Unicast Routing and VLT Multicast Routing.

Spanned VLANs

Any VLAN configured on both VLT peer nodes is referred to as a Spanned VLAN. The VLT Interconnect (VLTi) port is automatically added as a member of the Spanned VLAN. As a result, any adjacent router connected to at least one VLT node on a Spanned VLAN subnet is directly reachable from both VLT peer nodes at the routing level.

VLT Unicast Routing

VLT unicast routing is supported on the Z9000 S4810 S4820T platform.

VLT unicast routing locally routes packets destined for the L3 endpoint of the VLT peer. This method avoids suboptimal routing. Peer-routing syncs the MAC addresses of both VLT peers and requires two local DA entries in TCAM. In case a VLT node is down, a timer that allows you to configure the amount of time needed for peer recovery provides resiliency. You can enable VLT unicast across multiple configurations using VLT links. You can enable ECMP on VLT nodes using VLT unicast.

VLT unicast routing is supported on IPv4 only. To enable VLT unicast routing, both VLT peers must be in L3 mode. Static route and routing protocols such as RIP, OSPF, ISIS, and BGP are supported. However, point-to-point configuration is not supported. To enable VLT unicast, VLAN configuration must be symmetrical on both peers. You cannot configure the same VLAN as Layer 2 on one node and as Layer 3 on the other node. Configuration mismatches are logged in the syslog and display in the `show vlt mismatch` command output.

If you enable VLT unicast routing, the following actions occur:

- L3 routing is enabled on any new IP address configured for a VLAN interface that is up.

- L3 routing is enabled on any VLAN with an admin state of up.

NOTE: If the CAM is full, do not enable peer-routing.

Configuring VLT Unicast

To enable and configure VLT unicast, follow these steps.

1. Enable VLT on a switch, then configure a VLT domain and enter VLT-domain configuration mode.
   CONFIGURATION mode
   `vlt domain domain-id`
2. Enable peer-routing.
   VLT DOMAIN mode
   `peer-routing`
3. Configure the peer-routing timeout.
   VLT DOMAIN mode
   `peer-routing—timeout value`
   value: Specify a value (in seconds) from 1 to 65535.

VLT Multicast Routing

VLT multicast routing is supported on the Z9000 S4810 S4820T platform.

VLT Multicast Routing provides resiliency to multicast routed traffic during the multicast routing protocol convergence period after a VLT link or VLT peer fails using the least intrusive method (PIM) and does not alter current protocol behavior.

Unlike VLT Unicast Routing, a normal multicast routing protocol does not exchange multicast routes between VLT peers. When you enable VLT Multicast Routing, the multicast routing table is synced between the VLT peers. Only multicast routes configured with a Spanned VLAN IP as their IIF are synced between VLT peers. For multicast routes with a Spanned VLAN IIF, only OIFs configured with a Spanned VLAN IP interface are synced between VLT peers.

The advantages of syncing the multicast routes between VLT peers are:

- **VLT resiliency** — After a VLT link or peer failure, if the traffic hashes to the VLT peer, the traffic continues to be routed using multicast until the PIM protocol detects the failure and adjusts the multicast distribution tree.
- **Optimal routing** — The VLT peer that receives the incoming traffic can directly route traffic to all downstream routers connected on VLT ports.
- **Optimal VLTi forwarding** — Only one copy of the incoming multicast traffic is sent on the VLTi for routing or forwarding to any orphan ports, rather than forwarding all the routed copies.

Important Points to Remember

- You cannot configure a VLT node as a rendezvous point (RP), but any PIM-SM compatible VLT node can serve as a designated router (DR).
- You can only use one spanned VLAN from a PIM-enabled VLT node to an external neighboring PIM router.
- If you connect multiple spanned VLANs to a PIM neighbor, or if both spanned and non-spanned VLANs can access the PIM neighbor, ECMP can cause the PIM protocol running on each VLT peer node to choose a different VLAN or IP route to reach the PIM neighbor. This can result in issues with multicast route syncing between peers.
- Both VLT peers require symmetric Layer 2 and Layer 3 configurations on both VLT peers for any spanned VLAN.
- For optimal performance, configure the VLT VLAN routing metrics to prefer VLT VLAN interfaces over non-VLT VLAN interfaces.

- When using factory default settings on a new switch deployed as a VLT node, packet loss may occur due to the requirement that all ports must be open.
- ECMP is not compatible on VLT nodes using VLT multicast. You must use a single VLAN.

Configuring VLT Multicast

To enable and configure VLT multicast, follow these steps.

1. Enable VLT on a switch, then configure a VLT domain and enter VLT-domain configuration mode.

   CONFIGURATION mode

   `vlt domain` *`domain-id`*
2. Enable peer-routing.

   VLT DOMAIN mode

   `peer-routing`
3. Configure the multicast peer-routing timeout.

   VLT DOMAIN mode

   `multicast peer-routing—timeout` *`value`*

   `value`: Specify a value (in seconds) from 1 to 1200.
4. Configure a PIM-SM compatible VLT node as a designated router (DR). For more information, refer to Configuring a Designated Router.
5. Configure a PIM-enabled external neighboring router as a rendezvous point (RP). For more information, refer to Configuring a Static Rendezvous Point.
6. Configure the VLT VLAN routing metrics to prefer VLT VLAN interfaces over non-VLT VLAN interfaces. For more information, refer to Classify Traffic.
7. Configure symmetrical Layer 2 and Layer 3 configurations on both VLT peers for any spanned VLAN.

Non-VLT ARP Sync

Synchronization for non-ARP routing table entries is supported on the Z9000 S4810 S4820T platform.

Prior to FTOS version 9.2.(0.0), only ARP entries learned on VLT ports were synced between peers. In 9.2(0.0), ARP entries (including ND entries) learned on other ports are synced with the VLT peer to support station move scenarios.

NOTE: ARP entries learned on non-VLT, non-spanned VLANs are not synced with VLT peers.

Verifying a VLT Configuration

To monitor the operation or verify the configuration of a VLT domain, use any of the following `show` commands on the primary and secondary VLT switches.

- Display information on backup link operation.

  EXEC mode

  `show vlt backup-link`
- Display general status information about VLT domains currently configured on the switch.

  EXEC mode

  `show vlt brief`
- Display detailed information about the VLT-domain configuration, including local and peer port-channel IDs, local VLT switch status, and number of active VLANs on each port channel.

  EXEC mode

  `show vlt detail`

- Display the VLT peer status, role of the local VLT switch, VLT system MAC address and system priority, and the MAC address and priority of the locally-attached VLT device.

  EXEC mode

  `show vlt role`
- Display the current configuration of all VLT domains or a specified group on the switch.

  EXEC mode

  `show running-config vlt`
- Display statistics on VLT operation.

  EXEC mode

  `show vlt statistics`
- Display the RSTP configuration on a VLT peer switch, including the status of port channels used in the VLT interconnect trunk and to connect to access devices.

  EXEC mode

  `show spanning-tree rstp`
- Display the current status of a port or port-channel interface used in the VLT domain.

  EXEC mode

  `show interfaces` *`interface`*

  – *`interface`*: specify one of the following interface types:

    * Fast Ethernet: enter `fastethernet` *`slot/port`*.
    * 1-Gigabit Ethernet: enter `gigabitethernet` *`slot/port`*.
    * 10-Gigabit Ethernet: enter `tengigabitethernet` *`slot/port`*.
    * Port channel: enter `port-channel {1-128}`.

Example of the `show vlt backup-link` Command

```
FTOS_VLTpeer1# show vlt backup-link

VLT Backup Link
-----------------
Destination:                10.11.200.18
Peer HeartBeat status:      Up
HeartBeat Timer Interval:   1
HeartBeat Timeout:          3
UDP Port:                   34998
HeartBeat Messages Sent:    1026
HeartBeat Messages Received: 1025

FTOS_VLTpeer2# show vlt backup-link

VLT Backup Link
-----------------
Destination:                10.11.200.20
Peer HeartBeat status:      Up
HeartBeat Timer Interval:   1
HeartBeat Timeout:          3
UDP Port:                   34998
HeartBeat Messages Sent:    1030
HeartBeat Messages Received: 1014
```

Example of the `show vlt brief` Command

```
FTOS_VLTpeer1# show vlt brief
VLT Domain Brief
------------------
```

```
  Domain ID:                      1000
  Role:                           Secondary
  Role Priority:                  32768
  ICL Link Status:                Up
  HeartBeat Status:               Up
  VLT Peer Status:                Up
  Local Unit Id:                  0
  Version:                        5(1)
  Local System MAC address:       00:01:e8:8a:e9:70
  Remote System MAC address:      00:01:e8:8a:e7:e7
  Configured System MAC address: 00:0a:0a:01:01:0a
  Remote system version:          5(1)
  Delay-Restore timer:            90 seconds

FTOS_VLTpeer2# show vlt brief
VLT Domain Brief
------------------
Domain ID:                      1000
  Role:                         Primary
  Role Priority:                32768
  ICL Link Status:              Up
  HeartBeat Status:             Up
  VLT Peer Status:              Up
  Local Unit Id:                1
  Version:                      5(1)
  Local System MAC address:     00:01:e8:8a:e7:e7
  Remote System MAC address:    00:01:e8:8a:e9:70
  Configured System MAC address: 00:0a:0a:01:01:0a
  Remote system version:        5(1)
  Delay-Restore timer:          90 seconds
```

Example of the `show vlt detail` Command

```
FTOS_VLTpeer1# show vlt detail

Local LAG Id Peer LAG Id Local Status Peer Status Active VLANs
------------ ----------- ------------ ----------- -------------
100          100         UP            UP         10, 20, 30
127          2           UP            UP        20, 30

FTOS_VLTpeer2# show vlt detail

Local LAG Id Peer LAG Id Local Status Peer Status Active VLANs
------------ ----------- ------------ ----------- -------------
2            127         UP           UP          20, 30
100          100         UP           UP          10, 20, 30
```

Example of the `show vlt role` Command

```
FTOS_VLTpeer1# show vlt role

VLT Role
----------
VLT Role:                   Primary
System MAC address:         00:01:e8:8a:df:bc
System Role Priority:       32768
Local System MAC address:   00:01:e8:8a:df:bc
Local System Role Priority: 32768

FTOS_VLTpeer2# show vlt role

VLT Role
----------
VLT Role:                   Secondary
System MAC address:         00:01:e8:8a:df:bc
```

```
System Role Priority:       32768
Local System MAC address:   00:01:e8:8a:df:e6
Local System Role Priority: 32768
```

Example of the `show running-config vlt` Command

```
FTOS_VLTpeer1# show running-config vlt
!
vlt domain 30
  peer-link port-channel 60
  back-up destination 10.11.200.18

FTOS_VLTpeer2# show running-config vlt
!
vlt domain 30
  peer-link port-channel 60
  back-up destination 10.11.200.20
```

Example of the `show vlt statistics` Command

```
FTOS_VLTpeer1# show vlt statistics

VLT Statistics
----------------
HeartBeat Messages Sent:     987
HeartBeat Messages Received: 986
ICL Hello's Sent:            148
ICL Hello's Received:        98

FTOS_VLTpeer2# show vlt statistics

VLT Statistics
----------------
HeartBeat Messages Sent:     994
HeartBeat Messages Received: 978
ICL Hello's Sent:            89
ICL Hello's Received:        89
```

Example of the `show spanning-tree rstp` Command

The bold section displays the RSTP state of port channels in the VLT domain. Port channel 100 is used in the VLT interconnect trunk (VLTi) to connect to VLT peer2. Port channels 110, 111, and 120 are used to connect to access switches or servers (vlt).

```
FTOS_VLTpeer1# show spanning-tree rstp brief

Executing IEEE compatible Spanning Tree Protocol
Root ID Priority 0, Address 0001.e88a.dff8
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID Priority 4096, Address 0001.e88a.d656
Configured hello time 2, max age 20, forward delay 15

Interface                                    Designated
Name    PortID  Prio Cost  Sts  Cost   Bridge ID PortID
---------- -------- ---- ------- --------- ------- ------------------
Po 1    128.2   128 200000 DIS       800   4096   0001.e88a.d656 128.2
Po 3    128.4   128 200000 DIS       800   4096   0001.e88a.d656 128.4
Po 4    128.5   128 200000 DIS       800   4096   0001.e88a.d656 128.5
Po 100 128.101 128 800    FWD(VLTi) 800   0      0001.e88a.dff8 128.101
Po 110 128.111 128 00     FWD(vlt)  800   4096   0001.e88a.d656 128.111
Po 111 128.112 128 200000 DIS(vlt)  800   4096   0001.e88a.d656 128.112
Po 120 128.121 128 2000   FWD(vlt)  800   4096   0001.e88a.d656 128.121

FTOS_VLTpeer2# show spanning-tree rstp brief
```

```
Executing IEEE compatible Spanning Tree Protocol
Root ID Priority 0, Address 0001.e88a.dff8
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID Priority 0, Address 0001.e88a.dff8
We are the root
Configured hello time 2, max age 20, forward delay 15

Interface                                    Designated
Name    PortID  Prio Cost  Sts      Cost    Bridge ID PortID
--------- -------- ---- ------- -------- - ------- -------------
Po 1    128.2   128 200000 DIS      0   0   0001.e88a.dff8 128.2
Po 3    128.4   128 200000 DIS      0   0   0001.e88a.dff8 128.4
Po 4    128.5   128 200000 DIS      0   0   0001.e88a.dff8 128.5
Po 100 128.101 128 800    FWD(VLTi)0   0   0001.e88a.dff8 128.101
Po 110 128.111 128 00     FWD(vlt) 0   0   0001.e88a.dff8 128.111
Po 111 128.112 128 200000 DIS(vlt) 0   0   0001.e88a.dff8 128.112
Po 120 128.121 128 2000   FWD(vlt) 0   0   0001.e88a.dff8 128.121
```

Additional VLT Sample Configurations

To configure VLT, configure a backup link and interconnect trunk, create a VLT domain, configure a backup link and interconnect trunk, and connect the peer switches in a VLT domain to an attached access device (switch or server).

Review the following examples of VLT configurations.

Configuring Virtual Link Trunking (VLT Peer 1)

Enable VLT and create a VLT domain with a backup-link and interconnect trunk (VLTi).

```
FTOS_VLTpeer1(conf)#vlt domain 999
FTOS_VLTpeer1(conf-vlt-domain)#peer-link port-channel 100
FTOS_VLTpeer1(conf-vlt-domain)#back-up destination 10.11.206.35
FTOS_VLTpeer1(conf-vlt-domain)#exit
```

Configure the backup link.

```
FTOS_VLTpeer1(conf)#interface ManagementEthernet 0/0
FTOS_VLTpeer1(conf-if-ma-0/0)#ip address 10.11.206.23/
FTOS_VLTpeer1(conf-if-ma-0/0)#no shutdown
FTOS_VLTpeer1(conf-if-ma-0/0)#exit
```

Configure the VLT interconnect (VLTi).

```
FTOS_VLTpeer1(conf)#interface port-channel 100
FTOS_VLTpeer1(conf-if-po-100)#no ip address
FTOS_VLTpeer1(conf-if-po-100)#channel-member fortyGigE 0/56,60
FTOS_VLTpeer1(conf-if-po-100)#no shutdown
FTOS_VLTpeer1(conf-if-po-100)#exit
```

Configure the port channel to an attached device.

```
FTOS_VLTpeer1(conf)#interface port-channel 110
FTOS_VLTpeer1(conf-if-po-110)#no ip address
FTOS_VLTpeer1(conf-if-po-110)#switchport
FTOS_VLTpeer1(conf-if-po-110)#channel-member fortyGigE 0/52
FTOS_VLTpeer1(conf-if-po-110)#no shutdown
FTOS_VLTpeer1(conf-if-po-110)#vlt-peer-lag port-channel 110
FTOS_VLTpeer1(conf-if-po-110)#end
```

Verify that the port channels used in the VLT domain are assigned to the same VLAN.

```
FTOS_VLTpeer1# show vlan id 10
Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
```

```
Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - Hyperpull tagged

   NUM Status Description Q Ports
   10  Active             U Po110(Fo 0/52)
                          T Po100(Fo 0/56,60)
```

Configuring Virtual Link Trunking (VLT Peer 2)

Enable VLT and create a VLT domain with a backup-link VLT interconnect (VLTi).

```
FTOS_VLTpeer2(conf)#vlt domain 999
FTOS_VLTpeer2(conf-vlt-domain)#peer-link port-channel 100
FTOS_VLTpeer2(conf-vlt-domain)#back-up destination 10.11.206.23
FTOS_VLTpeer2(conf-vlt-domain)#exit
```

Configure the backup link.

```
FTOS_VLTpeer2(conf)#interface ManagementEthernet 0/0
FTOS_VLTpeer2(conf-if-ma-0/0)#ip address 10.11.206.35/
FTOS_VLTpeer2(conf-if-ma-0/0)#no shutdown
FTOS_VLTpeer2(conf-if-ma-0/0)#exit
```

Configure the VLT interconnect (VLTi).

```
FTOS_VLTpeer2(conf)#interface port-channel 100
FTOS_VLTpeer2(conf-if-po-100)#no ip address
FTOS_VLTpeer2(conf-if-po-100)#channel-member fortyGigE 0/46,50
FTOS_VLTpeer2(conf-if-po-100)#no shutdown
FTOS_VLTpeer2(conf-if-po-100)#exit
```

Configure the port channel to an attached device.

```
FTOS_VLTpeer2(conf)#interface port-channel 110
FTOS_VLTpeer2(conf-if-po-110)#no ip address
FTOS_VLTpeer2(conf-if-po-110)#switchport
FTOS_VLTpeer2(conf-if-po-110)#channel-member fortyGigE 0/48
FTOS_VLTpeer2(conf-if-po-110)#no shutdown
FTOS_VLTpeer2(conf-if-po-110)#vlt-peer-lag port-channel 110
FTOS_VLTpeer2(conf-if-po-110)#end
```

Verify that the port channels used in the VLT domain are assigned to the same VLAN.

```
FTOS_VLTpeer2# show vlan id 10
Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - Hyperpull tagged

   NUM Status Description Q Ports
   10  Active             U Po110(Fo 0/48)
                          T Po100(Fo 0/46,50)
```

Verifying a Port-Channel Connection to a VLT Domain (From an Attached Access Switch)

On an access device, verify the port-channel connection to a VLT domain.

```
FTOS_TORswitch(conf)# show running-config interface port-channel 11
!
```

```
interface Port-channel 11
no ip address
switchport
channel-member fortyGigE 1/18,22
no shutdown
```

Troubleshooting VLT

To help troubleshoot different VLT issues that may occur, use the following information.

NOTE: For information on VLT Failure mode timing and its impact, contact your Dell Networking representative.

Table 11. Troubleshooting VLT

Description	Behavior at Peer Up	Behavior During Run Time	Action to Take
Bandwidth monitoring	A syslog error message and an SNMP trap is generated when the VLTi bandwidth usage goes above the 80% threshold and when it drops below 80%.	A syslog error message and an SNMP trap is generated when the VLTi bandwidth usage goes above its threshold.	Depending on the traffic that is received, the traffic can be offloaded inVLTi.
Domain ID mismatch	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message and an SNMP trap are generated.	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message and an SNMP trap are generated.	Verify the domain ID matches on both VLT peers.
FTOS Version mismatch	A syslog error message is generated.	A syslog error message is generated.	Follow the correct upgrade procedure for the unit with the mismatched FTOS version.
Remote VLT port channel status	N/A	N/A	Use the `show vlt detail` and `show vlt brief` commands to view the VLT port channel status information.
Spanning tree mismatch at global level	All VLT port channels go down on both VLT peers. A syslog error message is generated.	No traffic is passed on the port channels. A one-time informational syslog message is generated.	During run time, a loop may occur as long as the mismatch lasts. To resolve, enable RSTP on both VLT peers.
Spanning tree mismatch at port level	A syslog error message is generated.	A one-time informational syslog message is generated.	Correct the spanning tree configuration on the ports.

Description	Behavior at Peer Up	Behavior During Run Time	Action to Take
System MAC mismatch	A syslog error message and an SNMP trap are generated.	A syslog error message and an SNMP trap are generated.	Verify that the unit ID of VLT peers is not the same on both units and that the MAC address is the same on both units.
Unit ID mismatch	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message is generated.	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message is generated.	Verify the unit ID is correct on both VLT peers. Unit ID numbers must be sequential on peer units; for example, if Peer 1 is unit ID "0", Peer 2 unit ID must be "1'.
Version ID mismatch	A syslog error message and an SNMP trap are generated.	A syslog error message and an SNMP trap are generated.	Verify the FTOS software versions on the VLT peers is compatible. For more information, refer to the *Release Notes* for this release.
VLT LAG ID is not configured on one VLT peer	A syslog error message is generated. The peer with the VLT configured remains active.	A syslog error message is generated. The peer with the VLT configured remains active.	Verify the VLT LAG ID is configured correctly on both VLT peers.
VLT LAG ID mismatch	The VLT port channel is brought down. A syslog error message is generated.	The VLT port channel is brought down. A syslog error message is generated.	Perform a mismatch check after the VLT peer is established.

FC Flex IO Modules

This part provides a generic, broad-level description of the operations, capabilities, and configuration commands of the Fiber Channel (FC) Flex IO module.

36

Understanding and Working of the FC Flex IO Modules

This chapter provides a generic, broad-level description of the operations and functionality of the Fiber Channel (FC) Flex IO module, and contains the following sections:

- FC Flex IO Modules Overview
- FC Flex IO Module Capabilities and Operations
- Guidelines for Working with FC Flex IO Modules
- Processing of Data Traffic
- Installing and Configuring the Switch
- Interconnectivity of FC Flex IO Modules with Cisco MDS Switches

FC Flex IO Modules Overview

The Fibre Channel (FC) Flex IO module is supported on Dell Networking MXL 10/40GbE Switch and Dell PowerEdge IO Aggregator (IOA). The MXL and IOA switch installed with the FC Flex IO module functions as a top-of-rack edge switch that supports converged enhanced ethernet (CEE) traffic — Fibre channel over Ethernet (FCoE) for storage, inter-process communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics. Although the MXL 10/40GbE Switch and the I/O Aggregator can act as a FIP snooping bridge (FSB) to provide FCoE transit switch capabilities, the salient and significant advantage of deploying the FC Flex IO module is to enable more streamlined and cohesive FCoE N-port identifier virtualization (NPIV) proxy gateway functionalities. The NPIV proxy gateway (NPG) provides FCoE-FC bridging behavior.

The FC Flex IO module offers a rich, comprehensive set of FCoE functionalities on the M1000e chassis by splitting the Ethernet and Fibre Channel (FC) traffic at the edge of the chassis. The FC switches that are connected directly to the FC Flex IO module provide Fibre Channel capabilities because the FC Flex IO module does not support full fabric functionalities. With the separation of Ethernet and FC packets performed at the edge of the chassis itself, you can use the MXL 10/40GBE Switch or the Aggregator that contains an FC Flex IO module to connect to a SAN environment without the need for a separate ToR switch to operate as NPIV proxy gateways. The MXL 10/40GBE Switch or the Aggregator can function in NPIV proxy gateway mode when an FC Flex IO module is present or in the FIP snooping bridge (FSB) mode when all the ports are Ethernet ports.

The FC Flex IO module uses the same baseboard hardware of the MXL 10/40GBE Switch or the Aggregator and the M1000 chassis. You can insert the FC Flex IO module into any of the optional module slots of the MXL 10/40GBE Switch and it provides four FC ports per module. If you insert only one FC Flex IO module, four ports are supported; if you insert two FC Flex IO modules, eight ports are supported.

By installing an FC Flex IO module, you can enable the MXL 10/40GbE Switch and I/O Aggregator to directly connect to an existing FC SAN network. The FC Flex IO module uses the existing slots on the MXL 10/40GbE Switch and I/O Aggregator and provides four or eight FC ports up to speed of 8 GbE per second. You can connect all of the FC ports to the same FC SAN fabric to yield FC bandwidth of up to 64 Gb. It is possible to connect some of the ports to a different FC SAN fabric to provide access to multiple fabric devices.

In a typical Fibre Channel storage network topology, separate network interface cards (NICs) and host bus adapters (HBAs) on each server (two each for redundancy purposes) are connected to LAN and SAN networks respectively. These deployments typically include a ToR SAN switch in addition to a ToR LAN switch. By employing converged network adapters (CNAs) that the FC Flex IO module supports, CNAs are used to transmit FCoE traffic from the server instead of separate NIC and HBA devices. In such a scenario, you can determine whether the FC or SAN packets and the Ethernet or LAN packets must be split within the chassis or by using a ToR switch to perform this splitting.

If you want to segregate the LAN and SAN traffic within the chassis, you can employ switches such as the Dell M8428-k Converged 10GbE Switch or FC-only switches such as the Dell M5424 switch module. You can also use the S5000 Switch as a ToR switch to separate the LAN and SAN traffic at the ToR. By using the FC Flex IO module, you can optimally and effectively split the LAN and SAN traffic at the edge of the blade chassis itself. You can deploy the FC Flex IO module can be deployed in the enterprise and data center switching networks to leverage and derive the advantages of a converged Ethernet network.

The FC Flex IO module is not an FCF switch, but it offers FCoE capabilities from the server to the MXL and I/O Aggregator switches, and native FC capability in the uplink direction to the SAN switches. Although the FC Flex IO module does not support all of the FCF characteristics, such as full-blown name services or zone parameters, it presents the most flexbile solution in interoperating with third-party switches that enable the splitting of LAN and SAN traffic. With the MXL 10/40GbE Switch and I/O Aggregator being well-established appliances in the switch domain, you can install the FC Flex IO module to enhance and increase the converged Ethernet network performance and behavior. With the FC Flex IO module, the MXL 10/40GbE Switch and I/O Aggregator provide thirty-two 1GbE or 10 GbE server-facing ports and the option to add two FC Flex IO modules that offer up to 8 8Gb Fibre Channel ports for uplink traffic in addition to the fixed two 40GbE ports on the MXL 10/40GbE Switch and I/O Aggregator.

You can configure one of the following upstream (fabric-facing) FC ports:

- Two 40GbE and eight 8Gb FC ports
- Four 40GbE and four 8Gb FC ports
- Two 40GbE, four 10GbE, and four 8Gb FC ports
- Two 40GbE, four 10GBASE-T, and four 8Gb FC ports

FC Flex IO Module Capabilities and Operations

The FC Flex IO module has the following characteristics:

- You can install one or two FC Flex IO modules on the MXL 10/40GbE Switch and I/O Aggregator. Each module supports four FC ports.
- Each port can operate in 2G, 4G, or 8G of Fibre Channel speed.
- All ports on an FC Flex IO module can function in NPIV mode that enables connectivity to FC switches or directors, and also to multiple SAN topologies.
- It automatically senses the current speed when the port link is up. Valid link speeds are 2 Gbps, 4 Gbps, and 8 Gbps.
- By default, the FC ports are configured to operate in N port mode to connect to an F port on an FC switch in a fabric. You can apply only one FCoE map on an FC port. An N-Port is a port on the node of an FC device and is called a node port.
- There is a maximum of 64 server fabric login (FLOGI) requests or fabric discovery (FDISC) requests per server MAC address before being forwarded by the FC Flex IO module to the FC core switch. Without user configuration, only 32 server login sessions are permitted for each server MAC address. To increase the total number of sessions to 64, use the `max sessions` command.

- A distance of up to 300 meters is supported at 8 Gbps for Fibre Channel traffic.
- Multiple domains are supported in an NPIV proxy gateway (NPG).
- You cannot configure the MXL and Aggregator switches in Stacking mode if the switches contain the FC Flex IO module. Similarly, FC Flex IO modules do not function when you insert them in to a stack of MXL/Aggregrator switches.
- If the switch does not contain FC Flex modules, you cannot create a stack and a log message displays stating that stacking is not supported unless the switches contain only FC Flex modules.

Guidelines for Working with FC Flex IO Modules

The following guidelines apply to the FC Flex IO module:

- All the ports of FC Flex IO modules operate in FC mode, and do not support Ethernet mode.
- FC Flex IO modules are not supported in the chassis management controller (CMC) GUI.
- The only supported FCoE functionality is NPIV proxy gateway. You must configure the other FCoE services, such as name server, zone server, and login server on an external FC switch.
- With the FC Flex IO module, the MXL 10/40GbE Switch continues to support bare metal provisioning (BMP) on any Ethernet port. BMP is not supported on FC ports. BMP improves accessibility to the MXL 10/40GbE Switch by automatically loading pre-defined configurations and boot images that are stored in file servers. You can use BMP on a single switch or on multiple switches.
- FC Flex IOM module is a field-replaceable unit (FRU). Its memory type is electrically erasable programmable read-only memory (EEPROM), which enables it to save manufacturing information, such as the serial number. It is hot-swappable, assuming that the module that is removed is replaced by the same type of module in that same slot.
- The FC FlexIO does not have persistent storage for any runtime configuration. All the persistent storage for runtime configuration is on the MXL and IOA baseboard.
- With both FC Flex IO modules present in the MXL or I/O Aggregator switches, the power supply requirement and maximum thermal output are the same as these parameters needed for the M1000 chassis.
- Each port on the FC Flex IO module contains status indicators to denote the link status and transmission activity. For traffic that is being transmitted, the port LED shows a blinking green light. The Link LED displays solid green when a proper link with the peer is established. If there is no connectivity, the LEDs are not lit
- The MXL and IOA switches continue to operate in FCoE Gateway mode even if connectivity to a TOR switch does not exist.
- The I/O Aggregrator examines whether the FC Flex IO module is inserted into the switch. When the FC Flex IO module is present during the boot process, the switch runs in FCoE NPIV gateway mode by default.
- When an FC Flex IO module is present in the IO Aggregrator, the software auto-configures the DCB settings on the ports that support DCB and does not retrieve these settings from the ToR switch.
- Active fabric manager (AFM) is compatible with FC Flex IO modules.
- All SNMP MIBs that are supported for MXL and IOA switches apply equally for FC Flex IO modules. The interface MIB indicates the FC interface when you install the FC flex IO module. The interface MIB statistical counters compute and display the FC interface metrics.

- When the Dell Networking OS sends FC frames (the initial FLOGI or FLOGO messages), or converts FLOGI to FDISC messages or processes any internally generated FC frames, the software computes and verifies the FC cyclic redundancy check (CRC) value before sending the frame to FC ports.
- Fabric worldwide name (WWN) verification is available for eight FC ports. Single-switching WWN capability is provided when the switch operates in NPIV mode.
- With FC Flex IO modules, you can connect the IOA in Simple MUX mode to a single fabric.
- With FC Flex IO modules on an IOA, the FC port speed is set to auto. The following parameters are automatically configured on the ENode facing and FC ports
- Description: SAN_FABRIC
- Fabric-id: 1002
- Fcoe-vlan: 1002
- Fc-map: 0x0efc00
- Fcf-priority: 128
- Fka-adv-period: 8000mSec
- Keepalive: enable
- Vlan priority: 3
- On an IOA, the FCoE virtual local area network (VLAN) is automatically configured.
- With FC Flex IO modules on an IOA, the following DCB maps are applied on all of the ENode facing ports.
- dcb-map: SAN_DCB_MAP
- priority-group 0 bandwidth 50 pfc off
- priority-group 1 bandwidth 50 pfc on
- priority-pgid 0 0 0 1 0 0 0 0
- On I/O Aggregrators, uplink failure detection (UFD) is disabled if FC Flex IO module is present to allow server ports to communicate with the FC fabric even when the Ethernet upstream ports are not operationally up.
- You must ensure that the NPIV functionality is enabled on the upstream switches that operate as FC switches or FCoE forwarders (FCF) before you connect the FC port of the MXL or I/O Aggregator to these upstream switches.
- While storage traffic traverses through FC Flex IO modules and the Ethernet uplink port-channel status changes (with DCB enabled on an adjacent switch), FCoE traffic is disrupted. This problem does not occur if Ethernet traffic is not involved and only FCoE traffic is transmitted. Also, if DCB on the ToR switch is disabled, traffic disruption does not occur.

Port Numbering for FC Flex IO Modules

Even-numbered ports are at the bottom of the I/O panel and for modules odd-numbered ports are at the top of the I/O panel. When installed in a PowerEdge M1000e Enclosure, the MXL 10/40GbE Switch and Aggregator ports are numbered 33 to 56 from the bottom to the top of the switch. The following port numbering convention applies to the FC Flex IO module:

- In expansion slot 0, the ports are numbered 41 to 44.
- In expansion slot 1, the ports are numbered 49 to 52.

Installing the Optics

The following optical ports are supported on the FC Flex IO module using one of the supported breakout cables:

- 4G or 8G Fibre Channel small form-factor pluggable plus (SFP+) optics module and LC connectors over a distance of 150 meters.
- 4G or 8G Fibre Channel SFP+ optics module and LC connectors over a distance of 4 km.

CAUTION:
Electrostatic discharge (ESD) damage can occur if the components are mishandled. Always wear an ESD-preventive wrist or heel ground strap when handling the FC Flex IO module and its components.

WARNING: When working with optical fibres, follow all the warning labels and always wear eye protection. Never look directly into the end of a terminated or unterminated fibre or connector as it may cause eye damage.

1. – Position the optic so it is in the correct position. The optic has a key that prevents it from being inserted incorrectly.
   – Insert the optic into the port until it gently snaps into place.

NOTE:

1. When you cable the ports, be sure not to interfere with the airflow from the small vent holes above and below the ports.

Processing of Data Traffic

The Dell Networking OS determines the module type that is plugged into the slot. Based on the module type, the software performs the appropriate tasks. The FC Flex IO module encapsulates and decapsulates the FCoE frames. Any non-FCoE or non-FIP traffic is directly switched by the module, and only FCoE frames are processed and transmitted out of the Ethernet network.

When the external device sends FCoE data frames to the switch that contains the FC Flex IO module, the destination MAC address represents one of the Ethernet MAC addresses assigned to FC ports. Based on the destination address, the FCoE header is removed from the incoming packet and the FC frame is transmitted out of the FC port. The flow control mechanism is performed using per-priority flow control to ensure that frame loss does not occur owing to congestion of frames.

Operation of the FIP Application

The NPIV proxy gateway terminates the FIP sessions and responses to FIP messages. The FIP packets are intercepted by the FC Flex IO module and sent to the Dell Networking OS for further analysis. The FIP application responds to the FIP VLAN discovery request from the host based on the configured FCoE VLANs. For every ENode and VN_Port that is logged in, the FIP application responds to keepalive messages for the virtual channel. If the FC link becomes inactive or a logging off of the switch occurs, the FIP engine sends clear virtual link (CVL) messages to the host. The FIP application also responds to solicited advertisements from the end-device. In addition, the FIP application periodically sends advertisement packets to the end-devices for each FCF that is part of the NPIV proxy gateway.

If FC Flex IO modules are installed, the I/O Aggregator does not perform FIP snooping because the FIP frames are terminated on the switch for NPIV operations. However, on MXL Switches, you can configure the switch to operate in FIP Snooping or NPIV mode.

If the MXL 10/40GbE Switch functions in the NPIV mode and you attempt to set the uplink port to be a FCF or a bridge port, a warning message displays and the setting is not saved. On the Aggregator, if the FC module is present, the uplink ports are not automatically set up as FCF or bridge ports. The FC Flex module cannot function as both an NPIV proxy gateway and a FIP snooping bridge at the same time.

Operation of the NPIV Proxy Gateway

The NPIV application on the FC Flex IO module manages the FC functionalities configured in Dell Networking OS. After the FC link comes up, the gateway sends the initial FLOGI request to the connected switch using the switch and port WWN methods. After a successful login, the NPIV gateway sends a notification to inform the CNA that the FCF available to log in. The source address of the FIP advertisement and FIP discovery advertisement response contain the MAC address of the FC Flex IO module port. Depending on the number of login sessions on a particular FCF, the NPIV gateway can load-balance the login sessions from ENodes.

The NPIV application performs the FLOGI to FDISC conversion and sends the new FC frame on the associated FC ports. After the external switch responds to the FLOGI request, the NPIV gateway establishes the NPIV session and send the frame to the FIP application. The FIP application establishes virtual links to convert FCoE FLOGI accept messages into FIP FLOGI accept messages. The corresponding ACL for the accept message is then applied. If a FIP timeout from ENode or VN_PORT occurs, the NPIV application performs the FC fabric logout to the external FC switch. The NPIV application manages the sessions between the FCoE and the FC domain.

Installing and Configuring the Switch

After you unpack the MXL 10/40GbE Switch, refer to the flow chart in the following figure for an overview of the steps you must follow to install the blade and perform the initial configuration.

Installing and Configuring Flowchart for FC Flex IO Modules

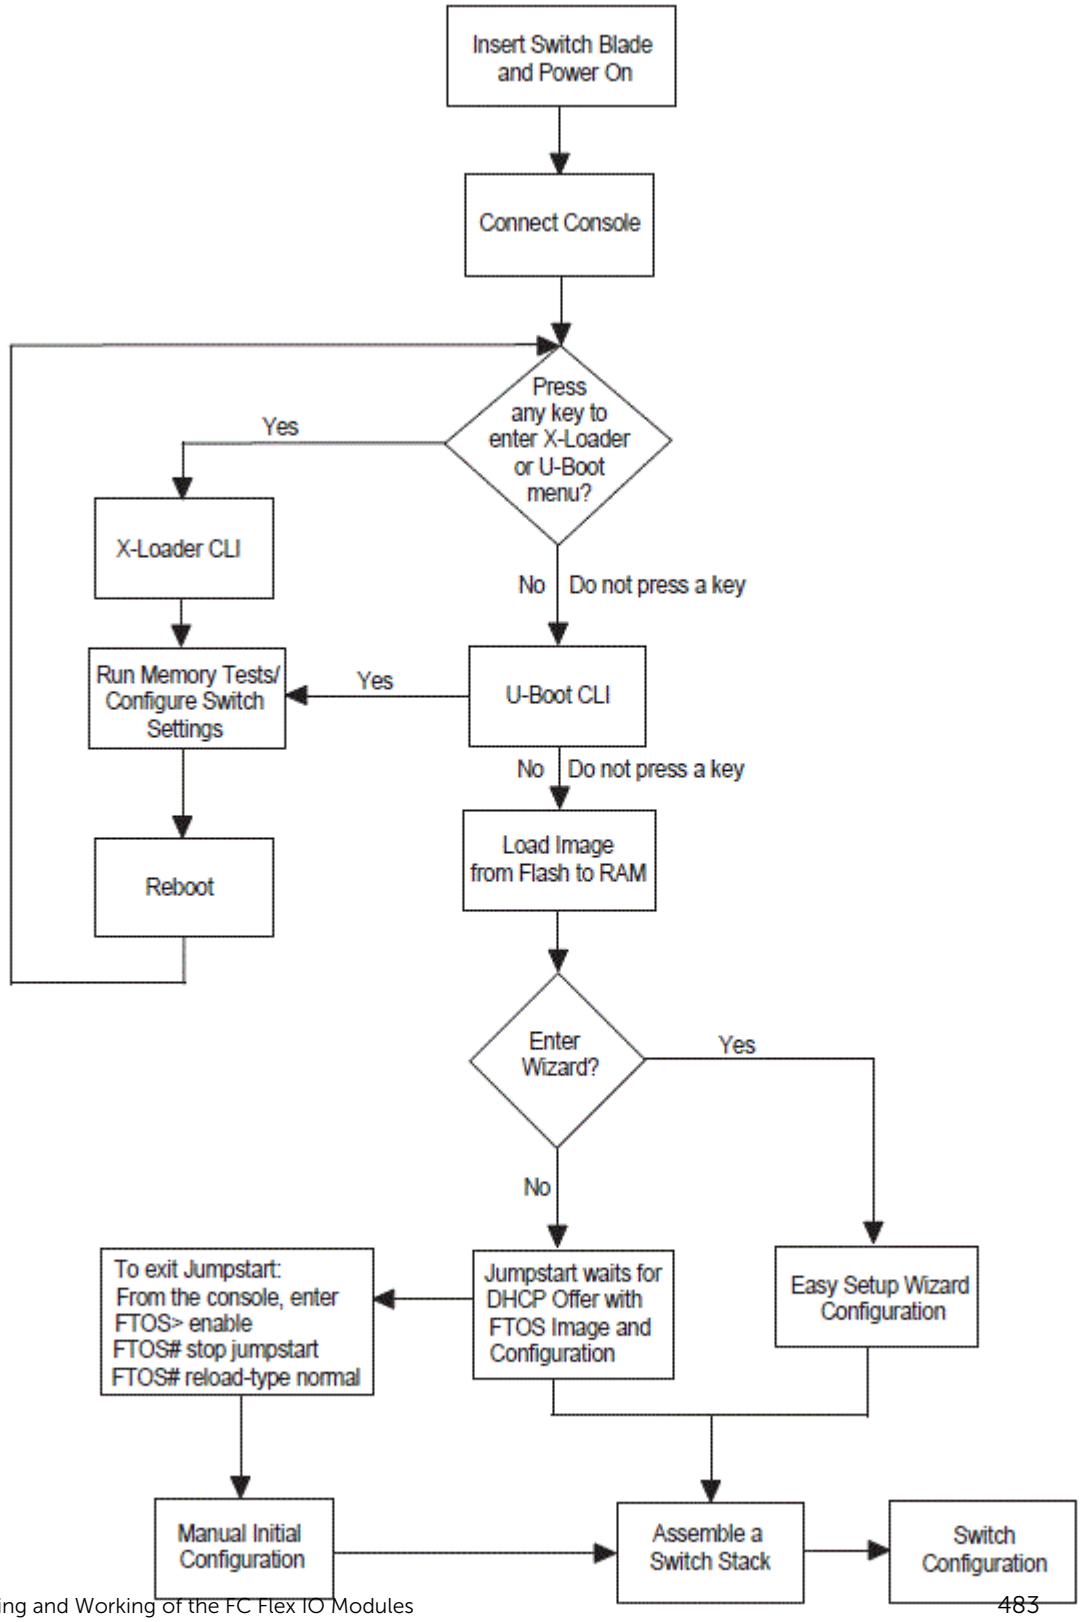

To see if a switch is running the latest Dell Networking OS version, use the **show version** command. To download a Dell Networking OS version, go to http://support.dell.com.

Installation

Site Preparation

Before installing the switch or switches, make sure that the chosen installation location meets the following site requirements:

- **Clearance** — There is adequate front and rear clearance for operator access. Allow clearance for cabling, power connections, and ventilation.
- **Cabling** — The cabling is routed to avoid sources of electrical noise such as radio transmitters, broadcast amplifiers, power lines, and fluorescent lighting fixtures.
- **Ambient Temperature** — The ambient switch operating temperature range is 10° to 35ºC (50° to 95ºF).

1. Decrease the maximum temperature by 1°C (1.8°F) per 300 m (985 ft.) above 900 m (2955 ft.).
2. **Relative Humidity** — The operating relative humidity is 8% to 85% (non-condensing) with a maximum humidity gradation of 10% per hour.

Unpacking the Switch

Package Contents

When unpacking each switch, make sure that the following items are included:

- One Dell Networking MXL 10/40GbE Switch IO Module
- One USB type A-to-DB-9 female cable
- Getting Started Guide
- Safety and Regulatory Information
- Warranty and Support Information
- Software License Agreement

Unpacking Steps

- Before unpacking the switch, inspect the container and immediately report any evidence of damage.
- Place the container on a clean, flat surface and cut all straps securing the container.
- Open the container or remove the container top.
- Carefully remove the switch from the container and place it on a secure and clean surface.
- Remove all packing material.
- Inspect the product and accessories for damage.

After you insert a FlexIO module into an empty slot, you must reload the Aggregator for the module. If you remove an installed module and insert a different module type, an error message displays to remind you that the slot is configured for a different type of FlexIO module. You must reload the switch to make the Flex IO module operational.

Interconnectivity of FC Flex IO Modules with Cisco MDS Switches

In a network topology that contains Cisco MDS switches, FC Flex IO modules that are plugged into the MXL and I/O Aggregator switches enable interoperation for a robust, effective deployment of the NPIV proxy gateway and FCoE-FC bridging behavior. In an environment that contains FC Flex IO modules and Cisco MDS switches, perform the following steps:

- Insert the FC Flex IO module into any of the optional module slots of the MXL 10/40GBE Switch or the I/O Aggregator Switch and reload the switch.
- When the device is reloaded, NPIV mode is automatically enabled.
- Configure the NPIV-related commands on MXL or I/O Aggregator.

After you perform the preceding procedure, the following operations take place:

- A physical link is established between the FC Flex I/O module and the Cisco MDS switch.
- The FC Flex I/O module sends a proxy FLOGI request to the upstream F_Port of the FC switch or the MDS switch. The F_port accepts the proxy FLOGI request for the FC FlexIO virtual N_Port. The converged network adapters (CNAs) are brought online and the FIP application is run.
- Discovery of the VLAN and FCF MAC addresses is completed.
- The CNA sends a FIP fabric login (FLOGI) request to the FC Flex IO module, which converts FLOGI to FDISC messages or processes any internally generated FC frames and sends these messages to the SAN environment.
- When the FC fabric discovery (FDISC) accept message is received from the SAN side, the Fc Flex IO module converts the FDISC message again into an FLOGI accept message and transmits it to the CNA.
- Internal tables of the switch are then programmed to enable the gateway device to forward FCoE traffic directly back and forth between the devices.
- The FC Flex IO module sends an FC or FCoE registered state change notification (RSCN) message to the upstream or downstream devices whenever an error occurs in the appropriate direction.
- An F_Port is a port on an FC switch that connects to an N_Port of an FC device and is called a fabric port.

By default, the NPIV functionality is disabled on the Cisco MDS switch; you must enable this capability before you connect the FC port of the MXL or I/O Aggregator to these upstream switches.

Data Center Bridging, Fibre Channel over Ethernet, and NPIV Proxy Gateway features are supported on the FC Flex IO modules. For detailed information about these applications and their working, see the corresponding chapters for these applications in this manual.

The following figures illustrate two deployment scenarios of configuring FC Flex IO modules:

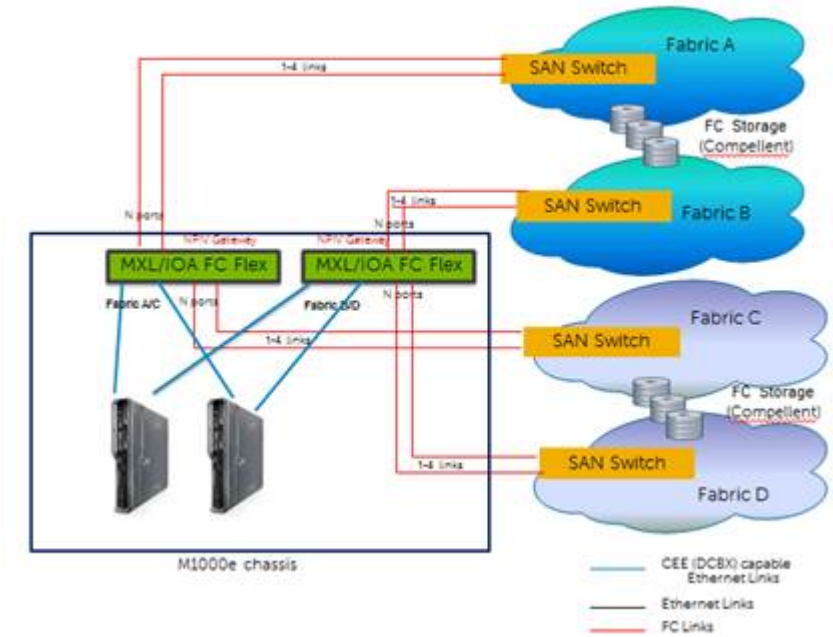

Figure 6. Case 1: Deployment Scenario of Configuring FC Flex IO Modules

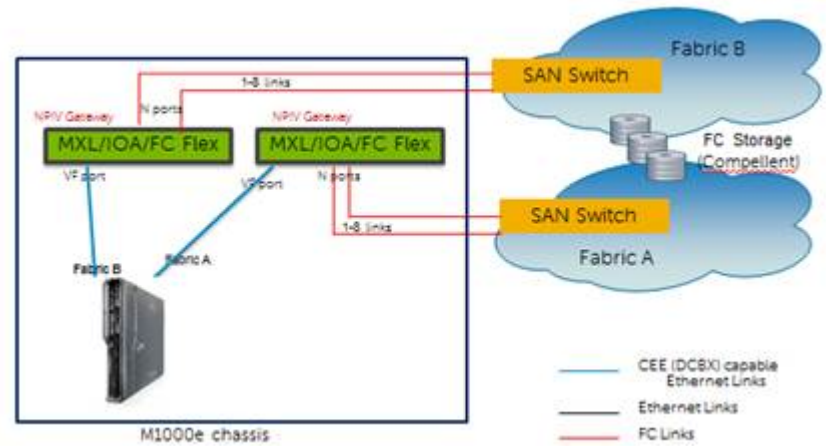

Figure 7. Case 2: Deployment Scenario of Configuring FC Flex IO Modules

37

Data Center Bridging (DCB) for FC Flex IO Modules

Data center bridging (DCB) refers to a set of IEEE Ethernet enhancements that provide data centers with a single, robust, converged network to support multiple traffic types, including local area network (LAN), server, and storage traffic.

The Fibre Channel (FC) Flex IO module is supported on Dell Networking MXL 10/40GbE Switch and Dell PowerEdge IO Aggregator (IOA). The MXL and IOA switch installed with the FC Flex IO module functions as a top-of-rack edge switch that supports converged enhanced ethernet (CEE) traffic — Fibre channel over Ethernet (FCoE) for storage, inter-process communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics.

The `dcb-input` and `dcb-output` configuration commands are deprecated, starting with Dell Networking OS Release 9.3.0.0 on the S4810, S6000, M I/O Aggregator, and MXL 10/40GbE Switch platforms. You must use the `dcp-map` command to create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic.

The Dell Networking operating software (Dell) commands for data center bridging features include 802.1Qbb priority-based flow control (PFC), 802.1Qaz enhanced transmission selection (ETS), and the data center bridging exchange (DCBX) protocol.

Interworking of DCB Map With DCB Buffer Threshold Settings

DCB map functionality is supported on the S4810, S4820T, S6000, I/O Aggregator, and MXL platforms.

The `dcb-input` and `dcb-output` configuration commands are deprecated. You must use the `dcp-map` command to create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic.

You can configure the `dcb-buffer-threshold` command and its related parameters only on ports with either auto configuration or dcb-map configuration. This command is not supported on existing front-panel interfaces or stack ports that are configured with the `dcb-input` or `dcb-output` commands. Similarly, if dcb-buffer-threshold configuration is present on any interface or a stack port, dcb-input or dcb-ouput policy cannot be applied on those interfaces.

Example: When dcb-buffer-threshold is applied on interfaces or stack ports with dcb-input or dcb-output policy, the following error message is displayed:

```
%Error: dcb-buffer-threshold not supported on interfaces with deprecated
commands
```

Example: When dcb-input or dcb-output is configured on interfaces or stack ports with dcb-buffer threshold policy:

```
%Error: Deprecated command is not supported on interfaces with dcb-buffer-
threshold configured
```

You must not modify the service-class dot1p mappings when any buffer-threshold-policy is configured on the system.

```
S4810-1(conf)#service-class dot1p-mapping dot1p0 3
% Error: PFC buffer-threshold policies conflict with dot1p mappings. Please
remove all dcb-buffer-threshold policies to change mappings.
```

The show dcb command has been enhanced to display the following additional buffer-related information:

```
 S4810-YU-MR-FTOS (conf)#do show dcb
 dcb Status :  Enabled
 PFC Queue Count :  2 --Indicate the PFC queue configured.
 Total buffer (lossy + lossless)(in KB): 7787--Total buffer space for lossy and
lossless queues
 PFC total buffer (in KB): 6526 --Indicates the total buffer (configured or
default)
 PFC shared buffer (in KB): 832--Indicates the shared buffer (Configured or
default)
 PFC available buffer ( in KB): 5694--Indicates remaining available buffers for
PFC that are free to be allocated
```

dcb-map

Create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic. Apply the DCB map to an Ethernet interface.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

```
dcb-map map-name
```

Parameters

map-name	Enter a DCB map name. The maximum number of alphanumeric characters is 32.

Defaults

None. There are no pre-configured PFC and ETS settings on S5000 Ethernet interfaces.

Command Modes

CONFIGURATION

INTERFACE

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

A DCB map is a template used to configure DCB parameters and apply them on converged Ethernet interfaces. DCB parameters include priority-based flow control (PFC) and enhanced traffic selection (ETS).

To display the PFC and ETS settings in DCB maps, enter the `show qos dcb-map` command.

Use the `dcb-map` command to create a DCB map to specify PFC and ETS settings and apply it on Ethernet ports. After you apply a DCB map to an interface, the PFC and ETS settings in the map are applied when the Ethernet port is enabled. DCBx is enabled on Ethernet ports by default.

The `dcb-map` command is supported only on physical Ethernet interfaces.

To remove a DCB map from an interface, enter the `no dcb-map map-name` command in Interface configuration mode.

Related Commands

show qos dcb-map– displays the dcb-map profiles configured on the system.

dcb-map stack-unit all stack-ports all– applies a DCB map on all ports of a switch stack.

priority-pgid

Assign 802.1p priority traffic to a priority group in a DCB map.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

`priority-pgid dot1p0_group-num dot1p1_group-num dot1p2_group-num dot1p3_group-num dot1p4_group-num dot1p5_group-num dot1p6_group-num dot1p7_group-num`

Parameters

dot1p0_group-num	Enter the priority group number for each 802.1p class of traffic in a DCB map.
dot1p1_group-num	
dot1p2_group-num	
dot1p3_group-num	
dot1p4_group-num	
dot1p5_group-num	
dot1p6_group-num	
dot1p7_group-num	

Defaults None

Command Modes DCB MAP

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

PFC and ETS settings are not pre-configured on Ethernet ports. You must use the `dcb-map` command to configure different groups of 802.1p priorities with PFC and ETS settings.

Using the `priority-pgid` command, you assign each 802.1p priority to one priority group. A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group. For example, the `priority-pgid 0 0 0 1 2 4 4 4` command creates the following groups of 802.1p priority traffic:

- Priority group 0 contains traffic with dot1p priorities 0, 1, and 2.
- Priority group 1 contains traffic with dot1p priority 3.
- Priority group 2 contains traffic with dot1p priority 4.
- Priority group 4 contains traffic with dot1p priority 5, 6, and 7.

To remove a priority-pgid configuration from a DCB map, enter the `no priority-pgid` command.

Related Commands

dcb-map — creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

priority-group bandwidth pfc— configures the ETS bandwidth allocation and the PFC setting used to manage the port traffic in an 802.1p priority group.

priority-group bandwidth pfc

Configure the ETS bandwidth allocation and PFC mode used to manage port traffic in an 802.1p priority group.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

`priority-group` *`group-num`* `{bandwidth` *`percentage`*`| strict-priority} pfc {on | off}`

Parameters

priority-group ***group-num***	Enter the keyword `priority-group` followed by the number of an 802.1p priority group. Use the `priority-pgid` command to create the priority groups in a DCB map.
bandwidth ***percentage***	Enter the keyword `bandwidth` followed by a bandwidth percentage allocated to the priority group. The range of valid values is 1 to 100. The sum of all allocated bandwidth percentages in priority groups in a DCB map must be 100%.
strict-priority	Configure the priority-group traffic to be handled with strict priority scheduling. Strict-priority traffic is serviced first, before bandwidth allocated to other priority groups is made available.
pfc {on | off}	Configure whether priority-based flow control is enabled (on) or disabled (off) for port traffic in the priority group.

Defaults None

Command Modes DCB MAP

Command History	**Version 9.3.0.0**	Introduced on the S4810 and S6000 platforms.
	Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

Use the `dcb-map` command to configure priority groups with PFC and/or ETS settings and apply them to Ethernet interfaces.

Use the `priority-pgid` command to map 802.1p priorities to a priority group. You can assign each 802.1p priority to only one priority group. A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.

Repeat the `priority-group bandwidth pfc` command to configure PFC and ETS traffic handling for each priority group in a DCB map.

You can enable PFC on a maximum of two priority queues.

If you configure more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.

To remove a priority-group configuration in a DCB map, enter the `no priority-group bandwidth pfc` command.

By default, equal bandwidth is assigned to each dot1p priority in a priority group. Use the `bandwidth` parameter to configure the bandwidth percentage assigned to a priority group. The sum of the bandwidth allocated to all priority groups in a DCB map must be 100% of the bandwidth on the link. You must allocate at least 1% of the total port bandwidth to each priority group.

Related Commands

dcb-map – creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

priority-pgid – configures the 802.1p priority traffic in a priority group for a DCB map.

dcb-map stack-unit all stack-ports all

Apply the specified DCB map on all ports of the switch stack.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax

`dcb-map stack-unit all stack-ports all` *`dcb-map-name`*

To remove the PFC and ETS settings in a DCB map from all stack units, use the `no dcb-map stack-unit all stack-ports all` command.

Parameters

dcb-map-name	Enter the name of the DCB map.

Defaults

none

Command Modes	CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information The `dcb-map stack-unit all stack-ports all` command overwrites any previous DCB maps applied to stack ports.

Related Commands dcb-map – creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

show qos dcb-map

Display the DCB parameters configured in a specified DCB map.

S4810, S6000, and FC Flex IO Modules with MXL and I/O Aggregator

Syntax `show qos dcb-map map-name`

Parameters

map-name	Displays the PFC and ETS parameters configured in the specified map.

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the S4810 and S6000 platforms.
Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information Use the `show qos dcb-map` command to display the enhanced transmission selection (ETS) and priority-based flow control (PFC) parameters used to configure server-facing Ethernet ports. S5000 Ethernet ports are DCBx-enabled by default.

The following table describes the `show qos dcb-map` output shown in the example below.

Field	Description
State	Complete: All mandatory DCB parameters are correctly configured. In progress: The DCB map configuration is not complete. Some mandatory parameters are not configured.
PFC Mode	PFC configuration in DCB map: On (enabled) or Off.
PG	Priority group configured in the DCB map.

Field	Description
TSA	Transmission scheduling algorithm used by the priority group: Enhanced Transmission Selection (ETS).
BW	Percentage of bandwidth allocated to the priority group.
PFC	PFC setting for the priority group: On (enabled) or Off.
Priorities	802.1p priorities configured in the priority group.

Example

```
FTOS# show qos dcb-map dcbmap2

State   :Complete
PfcMode:ON
--------------------
PG:0 TSA:ETS  BW:50  PFC:OFF
Priorities:0 1 2 4 5 6 7


PG:1 TSA:ETS  BW:50  PFC:ON
Priorities:3
```

Related Commands

dcb-map — creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

DCB Command

The following DCB command is supported on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch, and S4810 and S6000 platforms.

dcb-enable

Enable data center bridging.

Syntax

`dcb enable`

To disable DCB, use the `no dcb enable` command.

Defaults

none

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

DCB is not supported if you enable `link-level flow control` on one or more interfaces.

DCBX Commands

The following DCBX commands are supported on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch, and S4810 and S6000 platforms.

advertise dcbx-appln-tlv

On a DCBX port with a manual role, configure the application priority TLVs advertised on the interface to DCBX peers.

Syntax

```
advertise dcbx-appln-tlv {fcoe | iscsi}
```

To remove the application priority TLVs, use the `no advertise dcbx-appln-tlv {fcoe | iscsi}` command.

Parameters

{fcoe | iscsi} Enter the application priority TLVs, where:

- `fcoe`: enables the advertisement of FCoE in application priority TLVs.
- `iscsi`: enables the advertisement of iSCSI in application priority TLVs.

Defaults Application priority TLVs are enabled to advertise FCoE and iSCSI.

Command Modes PROTOCOL LLDP

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information To disable TLV transmission, use the `no` form of the command; for example, `no advertise dcbx-appln-tlv iscsi`.

advertise dcbx-tlv

On a DCBX port with a manual role, configure the PFC and ETS TLVs advertised to DCBX peers.

Syntax

```
advertise dcbx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]
```

To remove the advertised ETS TLVs, use the `no advertise dcbx-tlv` command.

Parameters

{ets-conf | ets-reco | pfc} Enter the PFC and ETS TLVs advertised, where:

- `ets-conf`: enables the advertisement of ETS configuration TLVs.
- `ets-reco`: enables the advertisement of ETS recommend TLVs.
- `pfc`: enables the advertisement of PFC TLVs.

Defaults All PFC and ETS TLVs are advertised.

Command Modes PROTOCOL LLDP

Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

You can configure the transmission of more than one TLV type at a time; for example: `advertise dcbx-tlv ets-conf ets-reco`.

You can enable ETS recommend TLVs (`ets-reco`) only if you enable ETS configuration TLVs (`ets-conf`). To disable TLV transmission, use the `no` form of the command; for example, `no advertise dcbx-tlv pfc ets-reco`.

DCBX requires that you enable LLDP to advertise DCBX TLVs to peers.

Configure DCBX operation at the INTERFACE level on a switch or globally on the switch. To verify the DCBX configuration on a port, use the `show interface dcbx detail` command.

dcbx port-role

Configure the DCBX port role the interface uses to exchange DCB information.

Syntax

`dcbx port-role {config-source | auto-downstream | auto-upstream | manual}`

To remove DCBX port role, use the `no dcbx port-role {config-source | auto-downstream | auto-upstream | manual}` command.

Parameters

config-source | auto-downstream | auto-upstream | manual

Enter the DCBX port role, where:

- `config-source`: configures the port to serve as the configuration source on the switch.
- `auto-upstream`: configures the port to receive a peer configuration. The configuration source is elected from auto-upstream ports.
- `auto-downstream`: configures the port to accept the internally propagated DCB configuration from a configuration source.
- `manual`: configures the port to operate only on administer-configured DCB parameters. The port does not accept a DCB configuration received form a peer or a local configuration source.

Defaults

Manual

Command Modes

INTERFACE PROTOCOL LLDP

Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

DCBX requires that you enable LLDP to advertise DCBX TLVs to peers.

Configure DCBX operation at the INTERFACE level on a switch or globally on the switch. To verify the DCBX configuration on a port, use the `show interface dcbx detail` command.

dcbx version

Configure the DCBX version used on the interface.

Syntax

`dcbx version {auto | cee | cin | ieee-v2.5}`

To remove the DCBX version, use the `dcbx version {auto | cee | cin | ieee-v2.5}` command.

Parameters

auto | cee | cin | ieee-v2.5

Enter the DCBX version type used on the interface, where:

- `auto`: configures the port to operate using the DCBX version received from a peer.
- `cee`: configures the port to use CDD (Intel 1.01).
- `cin`: configures the port to use Cisco-Intel-Nuova (DCBX 1.0).
- `ieee-v2`: configures the port to use IEEE 802.1az (Draft 2.5).

Defaults

Auto

Command Modes

INTERFACE PROTOCOL LLDP

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

DCBX requires that you enable LLDP to advertise DCBX TLVs to peers.

Configure DCBX operation at the INTERFACE level on a switch or globally on the switch. To verify the DCBX configuration on a port, use the `show interface dcbx detail` command.

debug dcbx

Enable DCBX debugging.

Syntax

`debug dcbx {all | auto-detect-timer | config-exchng | fail | mgmt | resource | sem | tlv}`

To disable DCBX debugging, use the `no debug dcbx` command.

Parameters	**{all | auto-detect-timer | config-exchng | fail | mgmt | resource | sem | tlv}**	Enter the type of debugging, where: • `all`: enables all DCBX debugging operations. • `auto-detect-timer`: enables traces for DCBX auto-detect timers. • `config-exchng`: enables traces for DCBX configuration exchanges. • `fail`: enables traces for DCBX failures. • `mgmt`: enables traces for DCBX management frames. • `resource`: enables traces for DCBX system resource frames. • `sem`: enables traces for the DCBX state machine. • `tlv`: enables traces for DCBX TLVs.
Defaults	none	
Command Modes	EXEC Privilege	
Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

fcoe priority-bits

Configure the FCoE priority advertised for the FCoE protocol in application priority TLVs.

Syntax	`fcoe priority-bits` *priority-bitmap* To remove the configured FCoE priority, use the `no fcoe priority-bits` command.	
Parameters	***priority-bitmap***	Enter the priority-bitmap range. The range is from 1 to FF.
Defaults	0x8	
Command Modes	PROTOCOL LLDP	
Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch
Usage Information	This command is available at the global level only.	

iscsi priority-bits

Configure the iSCSI priority advertised for the iSCSI protocol in application priority TLVs.

Syntax

`iscsi priority-bits` *priority-bitmap*

To remove the configured iSCSI priority, use the `no iscsi priority-bits` command.

Parameters

priority-bitmap	Enter the priority-bitmap range. The range is from 1 to FF.

Defaults

0x10

Command Modes

PROTOCOL LLDP

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

This command is available at the global level only.

show interface dcbx detail

Displays the DCBX configuration on an interface.

Syntax

`show interface` *port-type slot/port* `dcbx detail`

Parameters

port-type	Enter the port type.
slot/port	Enter the slot/port number.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

To clear DCBX frame counters, use the `clear dcbx counters interface` *stack-unit/port* command.

The following describes the `show interface dcbx detail` command shown in the following example.

Field	Description
Interface	Interface type with chassis slot and port number.
Port-Role	Configured the DCBX port role: auto-upstream, auto-downstream, config-source, or manual.

Field	Description
DCBX Operational Status	Operational status (enabled or disabled) used to elect a configuration source and internally propagate a DCB configuration. The DCBX operational status is the combination of PFC and ETS operational status.
Configuration Source	Specifies whether the port serves as the DCBX configuration source on the switch: true (yes) or false (no).
Local DCBX Compatibility mode	DCBX version accepted in a DCB configuration as compatible. In auto-upstream mode, a port can only receive a DCBX version supported on the remote peer.
Local DCBX Configured mode	DCBX version configured on the port: CEE, CIN, IEEE v2.5, or Auto (port auto-configures to use the DCBX version received from a peer).
Peer Operating version	DCBX version that the peer uses to exchange DCB parameters.
Local DCBX TLVs Transmitted	Transmission status (enabled or disabled) of advertised DCB TLVs (see TLV code at the top of the show command output).
Local DCBX Status: DCBX Operational Version	DCBX version advertised in Control TLVs.
Local DCBX Status: DCBX Max Version Supported	Highest DCBX version supported in Control TLVs.
Local DCBX Status: Sequence Number	Sequence number transmitted in Control TLVs.
Local DCBX Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs.
Local DCBX Status: Protocol State	Current operational state of the DCBX protocol: ACK or IN-SYNC.
Peer DCBX Status: DCBX Operational Version	DCBX version advertised in Control TLVs received from the peer device.
Peer DCBX Status: DCBX Max Version Supported	Highest DCBX version supported in Control TLVs received from the peer device.

Field	Description
Peer DCBX Status: Sequence Number	Sequence number transmitted in Control TLVs received from the peer device.
Peer DCBX Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs received from the peer device.
Total DCBX Frames transmitted	Number of DCBX frames sent from the local port.
Total DCBX Frames received	Number of DCBX frames received from the remote peer port.
Total DCBX Frame errors	Number of DCBX frames with errors received.
Total DCBX Frames unrecognized	Number of unrecognizable DCBX frames received.

Example

```
Dell(conf)# show interface tengigabitethernet 0/49 dcbx detail
Dell#show interface te 0/49 dcbx detail

E-ETS Configuration TLV enabled
  e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled
  r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled
  p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled
  f-Application Priority for FCOE disabled
I-Application priority for iSCSI enabled
  i-Application Priority for iSCSI disabled
---------------------------------------------------------

Interface TenGigabitEthernet 0/49
  Remote Mac Address 00:00:00:00:00:11
   Port Role is Auto-Upstream
  DCBX Operational Status is Enabled
  Is Configuration Source? TRUE

Local DCBX Compatibility mode is CEE
  Local DCBX Configured mode is CEE
  Peer Operating version is CEE
  Local DCBX TLVs Transmitted: ErPfi

Local DCBX Status
-----------------
  DCBX Operational Version is 0
  DCBX Max Version Supported is 0
  Sequence Number: 2
  Acknowledgment Number: 2
  Protocol State: In-Sync

Peer DCBX Status:
----------------
  DCBX Operational Version is 0
```

```
DCBX Max Version Supported is 255
Sequence Number: 2
Acknowledgment Number: 2
Total DCBX Frames transmitted 27
Total DCBX Frames received 6
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0
```

ETS Commands

The following ETS commands are supported on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch, and S4810 and S6000 platforms.

bandwidth-percentage

Configure the bandwidth percentage allocated to priority traffic in port queues.

Syntax

`bandwidth-percentage` *percentage*

To remove the configured bandwidth percentage, use the `no bandwidth-percentage` command.

Parameters

percentage	(Optional) Enter the bandwidth percentage. The percentage range is from 1 to 100% in units of 1%.

Defaults

none

Command Modes

QOS-POLICY-OUT-ETS

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

By default, equal bandwidth is assigned to each port queue and each dot1p priority in a priority group. To configure bandwidth amounts in associated dot1p queues, use the `bandwidth-percentage` command. When specified bandwidth is assigned to some port queues and not to others, the remaining bandwidth (100% minus assigned bandwidth amount) is equally distributed to unassigned nonstrict priority queues in the priority group. The sum of the allocated bandwidth to all queues in a priority group must be 100% of the bandwidth on the link.

ETS-assigned bandwidth allocation applies only to data queues, not to control queues.

The configuration of bandwidth allocation and strict-queue scheduling is not supported at the same time for a priority group. If you configure both, the configured bandwidth allocation is ignored for priority-group traffic when you apply the output policy on an interface.

By default, equal bandwidth is assigned to each priority group in the ETS output policy applied to an egress port if you did not configure bandwidth allocation. The sum of configured bandwidth allocation to dot1p priority traffic in all ETS priority groups must be 100%. Allocate at least 1% of the total bandwidth to each priority group and queue. If bandwidth is assigned to some priority groups but not to others, the remaining bandwidth (100% minus assigned bandwidth amount) is

equally distributed to nonstrict-priority groups which have no configured scheduler.

Related Commands

- qos-policy-output ets — creates a QoS output policy.
- scheduler — schedules priority traffic in port queues.

clear ets counters

Clear all ETS TLV counters on an interface.

Syntax `clear ets counters port-type slot/port`

Parameters

port-type	Enter the keywords `port-type` then the slot/port information.

Defaults none

Command Modes EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

dcb-map

Create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic. Apply the DCB map to an Ethernet interface.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax `dcb-map` *`map-name`*

Parameters

map-name	Enter a DCB map name. The maximum number of alphanumeric characters is 32.

Defaults None. There are no pre-configured PFC and ETS settings on M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module Ethernet interfaces. With auto-detection of DCB enabled, a DCB map named 'dcb-map' is applied on all the Ethernet interfaces on which the DCBx frames are observed.

Command Modes CONFIGURATION
INTERFACE

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information A DCB map is a template used to configure DCB parameters and apply them on converged Ethernet interfaces. DCB parameters include priority-based flow control (PFC) and enhanced traffic selection (ETS).

To display the PFC and ETS settings in DCB maps, enter the `show qos dcb-map` command.

Use the `dcb-map` command to create a DCB map to specify PFC and ETS settings and apply it on Ethernet ports. After you apply a DCB map to an interface, the PFC and ETS settings in the map are applied when the Ethernet port is enabled. DCBx is enabled on Ethernet ports by default.

The `dcb-map` command is supported only on physical Ethernet interfaces.

To remove a DCB map from an interface, enter the `no dcb-map map-name` command in Interface configuration mode.

Related Commands

show qos dcb-map– displays the dcb-map profiles configured on the system.

dcb-map stack-unit all stack-ports all– applies a DCB map on all ports of a switch stack.

dcb-output

To associate an ETS configuration with priority traffic, create a DCB output policy.

Syntax

`dcb-output policy-name`

To remove the ETS output policy globally, use the `no dcb output policy-name` command.

Parameters

policy-name	Enter the DCB output policy name. The maximum is 32 alphanumeric characters.

Defaults

none

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

To associate a priority group with an ETS output policy with scheduling and bandwidth configuration, create a DCB output policy. You can apply a DCB output policy on multiple egress ports. When you apply an ETS output policy on an interface, ETS-configured scheduling and bandwidth allocation take precedence over any configured settings in QoS output policies.

The ETS configuration associated with 802.1 priority traffic in a DCB output policy is used in DCBX negotiation with ETS peers.

Related Commands

dcb-policy output — applies the output policy.

dcb-policy output

Apply the output policy with the ETS configuration to an egress interface.

Syntax

`dcb-policy output` *`policy-name`*

To delete the output policy, use the `no dcb-policy output` command.

Parameters

policy-name	Enter the output policy name.

Defaults

none

Command Modes

INTERFACE

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

When you apply an ETS output policy to on interface, ETS-configured scheduling and bandwidth allocation take precedence over any configured settings in QoS output policies.

To remove an ETS output policy from an interface, use the `no dcb-policy output policy-name` command. ETS is enabled by default with the default ETS configuration applied (all dot1p priorities in the same group with equal bandwidth allocation).

Related Commands

dcb-output — creates a DCB output policy.

dcb-policy output stack-unit stack-ports all

Apply the specified DCB output policy on all ports of the switch stack or a single stacked switch.

Syntax

`dcb-policy output stack-unit {all | stack-unit-id} stack-ports all dcb-output-policy-name`

To remove all DCB output policies applied to the stacked ports, use the `no dcb-policy output stack-unit all` command.

To remove only the DCB output policies applied to the specified switch, use the `no dcb-policy output stack-unit` command.

Parameters

stack-unit-id	Enter the stack unit identification.
dcb-output-policy-name	Enter the policy name for the DCB output policy.

Defaults

none

Command Modes

CONFIGURATION

Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

The `dcb-policy output stack-unit all` command overwrites any previous `dcb-policy output stack-unit stack-unit-id` configurations. Similarly, a `dcb-policy output stack-unit stack-unit-id` command overwrites any previous `dcb-policy output stack-unit all` configuration.

You can apply a DCB output policy with ETS configuration to all stacked ports in a switch stack or an individual stacked switch. You can apply different DCB output policies to different stack units.

Related Commands

dcb-policy input stack-unit stack-ports all — applies the specified DCB input policy.

description

Enter a text description of the DCB policy (PFC input or ETS output).

Syntax

`description` *`text`*

To remove the text description, use the `no description` command.

Parameters

text	Enter the description of the output policy. The maximum is 32 characters.

Defaults

none

Command Modes

- DCB INPUT POLICY
- DCB OUTPUT POLICY

Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Related Commands

- dcb-input — creates a DCB PFC input policy.
- dcb-policy input — applies the output policy.
- dcb-output — creates a DCBETS output policy.
- dcb-policy output — applies the output policy.

ets mode on

Enable the ETS configuration so that scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBX TLV from a peer can take effect on an interface.

Syntax

`ets mode on`

To remove the ETS configuration, use the `no ets mode on` command.

Defaults	ETS mode is on.
Command Modes	DCB OUTPUT POLICY
Command History	**Version 9.3.0.0** Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch
Usage Information	If you disable ETS in an output policy applied to an interface using the `no ets mode on` command, any previously configured QoS settings at the interface or global level takes effect. If you configure QoS settings at the interface or global level and in an output policy map (the `service-policy output` command), the QoS configuration in the output policy takes precedence.
Related Commands	• dcb-output — creates a DCB output policy. • dcb-policy output — applies the output policy.

priority-group

To use with an ETS output policy, create an ETS priority group.

Syntax	`priority-group` *`group-name`* To remove the priority group, use the `no priority-group` command.
Parameters	***group-name*** Enter the name of the ETS priority group. The maximum is 32 characters.
Defaults	none
Command Modes	CONFIGURATION
Command History	**Version 9.3.0.0** Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.
Usage Information	A priority group consists of 802.1p priority values that are grouped for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group. You must configure 802.1p priorities in priority groups associated with an ETS output policy. You can assign each dot1p priority to only one priority group. The maximum number of priority groups supported in ETS output policies on an interface is equal to the number of data queues (4) on the port. The 802.1p priorities in a priority group can map to multiple queues. If you configure more than one priority queue as strict priority or more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

Related Commands	• priority-list — configures the 802.1p priorities for an ETS output policy. • set-pgid — configures the priority-group.

priority-group bandwidth pfc

Configure the ETS bandwidth allocation and PFC mode used to manage port traffic in an 802.1p priority group.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax

```
priority-group group-num {bandwidth percentage| strict-priority} pfc {on | off}
```

Parameters

priority-group *group-num*	Enter the keyword `priority-group` followed by the number of an 802.1p priority group. Use the `priority-pgid` command to create the priority groups in a DCB map.
bandwidth *percentage*	Enter the keyword `bandwidth` followed by a bandwidth percentage allocated to the priority group. The range of valid values is 1 to 100. The sum of all allocated bandwidth percentages in priority groups in a DCB map must be 100%.
strict-priority	Configure the priority-group traffic to be handled with strict priority scheduling. Strict-priority traffic is serviced first, before bandwidth allocated to other priority groups is made available.
pfc {on | off}	Configure whether priority-based flow control is enabled (on) or disabled (off) for port traffic in the priority group.

Defaults None

Command Modes DCB MAP

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

Use the `dcb-map` command to configure priority groups with PFC and/or ETS settings and apply them to Ethernet interfaces.

Use the `priority-pgid` command to map 802.1p priorities to a priority group. You can assign each 802.1p priority to only one priority group. A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.

Repeat the `priority-group bandwidth pfc` command to configure PFC and ETS traffic handling for each priority group in a DCB map.

You can enable PFC on a maximum of two priority queues.

If you configure more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.

To remove a priority-group configuration in a DCB map, enter the `no priority-group bandwidth pfc` command.

By default, equal bandwidth is assigned to each dot1p priority in a priority group. Use the `bandwidth` parameter to configure the bandwidth percentage assigned to a priority group. The sum of the bandwidth allocated to all priority groups in a DCB map must be 100% of the bandwidth on the link. You must allocate at least 1% of the total port bandwidth to each priority group.

Related Commands

dcb-map – creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

priority-pgid – configures the 802.1p priority traffic in a priority group for a DCB map.

priority-group qos-policy

Associate the 802.1p priority traffic in a priority group with the ETS configuration in a QoS output policy.

Syntax

`priority-group` *`group-name`* `qos-policy` *`ets-policy-name`*

To remove the 802.1p priority group, use the `no priority-group qos-policy` command.

Parameters

group-name	Enter the group name of the 802.1p priority group. The maximum is 32 characters.
ets-policy-name	Enter the ETS policy name.

Defaults

none

Command Modes

DCB OUTPUT POLICY

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

The ETS configuration associated with 802.1p priority traffic in a DCB output policy is used in DCBX negotiation with ETS peers.

If you disable ETS in an output policy applied to an interface using the `no ets mode on` command, any previously configured QoS settings at the interface or global level take effect. If you configure QoS settings at the interface or global level and in an output policy map (the `service-policy output` command), the QoS configuration in the output policy takes precedence.

Related Commands

- dcb-output — creates a DCB output policy.
- dcb-policy output — applies the output policy.

priority-list

Configure the 802.1p priorities for the traffic on which you want to apply an ETS output policy.

Syntax	`priority-list value` To remove the priority list, use the `no priority-list` command.	
Parameters	***value***	Enter the priority list value. Separate priority values with a comma; specify a priority range with a dash; for example, priority-list 3,5-7. The range is from 0 to 7.
Defaults	none	
Command Modes	PRIORITY-GROUP	
Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

By default:

- All 802.1p priorities are grouped in priority group 0.
- 100% of the port bandwidth is assigned to priority group 0. The complete bandwidth is equally assigned to each priority class so that each class has 12 to 13%.

Related Commands

- priority-group qos-policy — associates an ETS priority group with an ETS output policy.
- set-pgid — configures the priority-group.

qos-policy-output ets

To configure the ETS bandwidth allocation and scheduling for priority traffic, create a QoS output policy.

Syntax	`qos-policy-output policy-name ets` To remove the QoS output policy, use the `no qos-policy-output ets` command.	
Parameters	***policy-name***	Enter the policy name. The maximum is 32 characters.
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

If an error occurs in an ETS output-policy configuration, the configuration is ignored and the scheduler and bandwidth allocation settings are reset to the ETS default values (all priorities are in the same ETS priority group and bandwidth is allocated equally to each priority).

If an error occurs when a port receives a peer's ETS configuration, the port's configuration is reset to the previously configured ETS output policy. If no ETS output policy was previously applied, the port is reset to the default ETS parameters.

Related Commands

- scheduler — schedules the priority traffic in port queues.
- bandwidth-percentage — bandwidth percentage allocated to the priority traffic in port queues.

scheduler

Configure the method used to schedule priority traffic in port queues.

Syntax

`scheduler` *value*

To remove the configured priority schedule, use the `no scheduler` command.

Parameters

value	Enter schedule priority value. The valid values are: • `strict`: strict-priority traffic is serviced before any other queued traffic. • `werr`: weighted elastic round robin (werr) provides low-latency scheduling for priority traffic on port queues.

Defaults

Weighted elastic round robin (WERR) scheduling is used to queue priority traffic.

Command Modes

POLICY-MAP-OUT-ETS

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

dot1p priority traffic on the switch is scheduled to the current queue mapping. dot1p priorities within the same queue must have the same traffic properties and scheduling method.

ETS-assigned scheduling applies only to data queues, not to control queues.

The configuration of bandwidth allocation and strict-queue scheduling is not supported at the same time for a priority group. If you configure both, the configured bandwidth allocation is ignored for priority-group traffic when you apply the output policy on an interface.

Related Commands

- qos-policy-output ets — configures the ETS bandwidth allocation.

- bandwidth-percentage — bandwidth percentage allocated to priority traffic in port queues.

set-pgid

Configure the priority-group identifier.

Syntax	`set-pgid value` To remove the priority group, use the `no set-pgid` command.	
Parameters	***value***	Enter the priority group identification. The range is from 0 to 7.
Defaults	none	
Command Modes	PRIORITY-GROUP	
Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.
Related Commands	• priority-group qos-policy — creates an ETS priority group. • priority-list — configures the 802.1p priorities.	

show interface ets

Displays the ETS configuration applied to egress traffic on an interface, including priority groups with priorities and bandwidth allocation.

Syntax	`show interface port-type slot/port ets {summary | detail}`	
Parameters	***port-type slot/ port* ets**	Enter the port-type slot and port ETS information.
	{summary | detail}	Enter the keyword `summary` for a summary list of results or enter the keyword `detail` for a full list of results.
Command Modes	CONFIGURATION	
Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.
Usage Information	To clear ETS TLV counters, use the `clear ets counters interface port-type slot/port` command. The following describes the `show interface summary` command shown in the following example.	

Field	Description
Interface	Interface type with stack-unit and port number.
Max Supported TC Group	Maximum number of priority groups supported.
Number of Traffic Classes	Number of 802.1p priorities currently configured.
Admin mode	ETS mode: on or off. When on, the scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBX TLV from a peer can take effect on an interface.
Admin Parameters	ETS configuration on local port, including priority groups, assigned dot1p priorities, and bandwidth allocation.
Remote Parameters	ETS configuration on remote peer port, including admin mode (enabled if a valid TLV was received or disabled), priority groups, assigned dot1p priorities, and bandwidth allocation. If ETS admin mode is enabled on the remote port for DCBX exchange, the Willing bit received in ETS TLVs from the remote peer is included.
Local Parameters	ETS configuration on local port, including admin mode (enabled when a valid TLV is received from a peer), priority groups, assigned dot1p priorities, and bandwidth allocation.
Operational status (local port)	Port state for current operational ETS configuration: • `Init`: Local ETS configuration parameters were exchanged with the peer. • `Recommend`: Remote ETS configuration parameters were received from the peer. • `Internally propagated`: ETS configuration parameters were received from the configuration source.
ETS DCBX Oper status	Operational status of the ETS configuration on the local port: match or mismatch.
State Machine Type	Type of state machine used for DCBX exchanges of ETS parameters: Feature — for legacy DCBX versions; Asymmetric — for an IEEE version.
Conf TLV Tx Status	Status of ETS Configuration TLV advertisements: enabled or disabled.
Reco TLV Tx Status	Status of ETS Recommendation TLV advertisements: enabled or disabled.
Input Conf TLV pkts	Number of ETS Configuration TLVs received.
Output Conf TLV pkts	Number of ETS Configuration TLVs transmitted.

Field	Description
Error Conf TLV pkts	Number of ETS Error Configuration TLVs received.
Input Reco TLV pkts	Number of ETS Recommendation TLVs received.
Output Reco TLV pkts	Number of ETS Recommendation TLVs transmitted.
Error Reco TLV pkts	Number of ETS Error Recommendation TLVs received.

Example (Summary)

```
Dell(conf)# show interfaces te 0/0 ets summary
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters:
------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
0      0,1,2,3,4,5,6,7  100%      ETS
1                         0%      ETS
2                         0%      ETS
3                         0%      ETS
4                         0%      ETS
5                         0%      ETS
6                         0%      ETS
7                         0%      ETS
Priority#              Bandwidth  TSA
0                        13%      ETS
1                        13%      ETS
2                        13%      ETS
3                        13%      ETS
4                        12%      ETS
5                        12%      ETS
6                        12%      ETS
7                        12%      ETS
Remote Parameters:
-------------------
Remote is disabled
Local Parameters:
------------------
Local is enabled
TC-grp Priority#        Bandwidth TSA
0      0,1,2,3,4,5,6,7  100%      ETS
1                         0%      ETS
2                         0%      ETS
3                         0%      ETS
4                         0%      ETS
5                         0%      ETS
6                         0%      ETS
7                         0%      ETS
Priority#              Bandwidth TSA
0                        13%      ETS
1                        13%      ETS
2                        13%      ETS
3                        13%      ETS
4                        12%      ETS
```

```
5                        12%      ETS
6                        12%      ETS
7                        12%      ETS
Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
```

Example (Detail)

```
Dell(conf)# show interfaces tengigabitethernet 0/0 ets detail
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :
------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
0      0,1,2,3,4,5,6,7  100%      ETS
1                       0%        ETS
2                       0%        ETS
3                       0%        ETS
4                       0%        ETS
5                       0%        ETS
6                       0%        ETS
7                       0%        ETS

Priority#               Bandwidth TSA
0                       13%       ETS
1                       13%       ETS
2                       13%       ETS
3                       13%       ETS
4                       12%       ETS
5                       12%       ETS
6                       12%       ETS
7                       12%       ETS
Remote Parameters:
-------------------
Remote is disabled

Local Parameters :
------------------
Local is enabled
TC-grp Priority#        Bandwidth TSA
0      0,1,2,3,4,5,6,7  100%      ETS
1                       0%        ETS
2                       0%        ETS
3                       0%        ETS
4                       0%        ETS
5                       0%        ETS
6                       0%        ETS
7                       0%        ETS

Priority#               Bandwidth TSA
0                       13%       ETS
1                       13%       ETS
2                       13%       ETS
3                       13%       ETS
4                       12%       ETS
5                       12%       ETS
6                       12%       ETS
7                       12%       ETS
Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
```

```
0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf
TLV Pkts
0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV
Pkts, 0 Error Traffic Class
TLV
Pkts
```

show qos dcb-output

Displays the ETS configuration in a DCB output policy.

Syntax

```
show qos dcb-output [ets-profile]
```

Parameters

[ets-profile]	Enter the ETS profile.

Command Modes

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Example

```
Dell# show qos dcb-output
dcb-output ets
  priority-group san qos-policy san
  priority-group ipc qos-policy ipc
  priority-group lan qos-policy lan
```

show qos priority-groups

Displays the ETS priority groups configured on the switch, including the 802.1p priority classes and ID of each group.

Syntax

```
show qos priority-groups
```

Command Modes

EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Example

```
Dell#show qos priority-groups
priority-group ipc
  priority-list 4
  set-pgid 2
```

show stack-unit stack-ports ets details

Displays the ETS configuration applied to egress traffic on stacked ports, including ETS Operational mode on each unit and the configurated priority groups with dot1p priorities, bandwidth allocation, and scheduler type.

Syntax

```
show stack-unit {all | stack-unit} stack-ports {all | port-number} ets details
```

Parameters

stack-unit	Enter the stack unit identification.
port-number	Enter the port number.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Example

```
Dell(conf)# show stack-unit all stack-ports all ets details

Stack unit 0 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on

Admin Parameters:
--------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
-----------------------------------------------
0      0,1,2,3,4,5,6,7  100%      ETS
1                       - -
2                       - -
3                       - -
4                       - -
5                       - -
6                       - -
7                       - -
8                       - -

Stack unit 1 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on
Admin Parameters:
--------------------
Admin is enabled
TC-grp Priority#        Bandwidth TSA
-----------------------------------------------
0      0,1,2,3,4,5,6,7  100%      ETS
1                       -         -
2                       -         -
3                       -         -
4                       -         -
5                       -         -
6                       -         -
```

7	–	–
8	–	–

PFC Commands

The following PFC commands are supported on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch, and S4810 and S6000 platforms.

clear pfc counters

Clear the PFC TLV counters and PFC statistics on an interface or stack unit.

Syntax

```
clear pfc counters [port-type slot/port | stack-unit {unit number | all } all stack-ports all}] | interface {statistics}]
```

Parameters

port-type	Enter the keywords `port-type` then the slot/port information.
stack-unit *unit number*	Enter the keywords `stack-unit` then the stack-unit number to be cleared.
all stack-ports all	Enter the keywords `all stack-ports all` to clear the counters on all interfaces.
statistics	Enter the keyword `statistics` to clear only the hardware PFC counters.

Defaults none

Command Modes EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information If you do not use the `statistics` parameter, both hardware and DCBx counters clear.

dcb stack-unit pfc-buffering pfc-port-count pfc-queues

Configure the PFC buffer for all port pipes in a specified stack unit by specifying the port-pipe number, number of PFC-enabled ports, and number of configured lossless queues.

Syntax

```
dcb stack-unit stack-unit-id [port-set port-set-id] pfc-buffering pfc-ports {1-56} pfc-queues {1-2}
```

To remove the configuration for the PFC buffer on all port pipes in a specified stack unit, use the `no dcb stack-unit stack-unit-id [port-set port-set-id] pfc-buffering pfc-ports pfc-queues` command.

Parameters

stack-unit-id	Enter the stack unit identification. The range is from 0 to 5.

	port-set	Enter the port-set identification. The only valid port-set ID (port-pipe number) on an MXL Switch is 0.
	pfc-ports *{1-56}*	Enter the pfc-ports. The range is from 1 to 56.
	pfc-queues *{1-2}*	Enter the pfc-queue number. The range is from 1 to 2.

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

If you configure PFC on a 40GbE port, count the 40GbE port as four PFC-enabled ports in the pfc-port number you enter in the command syntax.

To achieve lossless PFC operation, the PFC port count and queue number used for the reserved buffer size that is created must be greater than or equal to the buffer size required for PFC-enabled ports and lossless queues on the switch.

You must reload the stack or a specified stack unit (use the `reload` command in EXEC Privilege mode) for the PFC buffer configuration to take effect.

Related Commands dcb stack-unit pfc-buffering pfc-port pfc-queues — configures the PFC buffer for all switches in the stack.

dcb-input

To apply pause or flow control for specified priorities using a configure delay time, create a DCB input policy.

Syntax `dcb-input` *`policy-name`*

To delete the DCB input policy, use the `no dcb-input` command.

Parameters

policy-name	Maximum: 32 alphanumeric characters.

Defaults none

Command Modes CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

As soon as you apply a DCB policy with PFC enabled on an interface, DCBx starts exchanging information with PFC-enabled peers. The IEEE802.1Qbb, CEE, and CIN versions of PFC TLV are supported. DCBx also validates PFC configurations received in TLVs from peer devices.

By applying a DCB input policy with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, also enable

PFC on all DCB egress ports or configure the dot1p priority-queue assignment of PFC priorities to lossless queues (refer to `pfc no-drop queues`).

To remove a DCB input policy, including the PFC configuration it contains, enter the `no dcb-input policy-name` command in Interface Configuration mode.

Related Commands dcb-policy input — applies the input policy with the PFC configuration.

dcb-policy input

Apply the input policy with the PFC configuration to an ingress interface.

Syntax `dcb-policy input policy-name`

To delete the input policy, use the `no dcb-policy input` command.

Parameters

policy-name	Enter the input policy name with the PFC configuration to an ingress interface.

Defaults none

Command Modes INTERFACE

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Usage Information

If you apply an input policy with PFC disabled (`no pfc mode on`):

- You can enable link-level flow control on the interface. To delete the input policy, first disable link-level flow control. PFC is then automatically enabled on the interface because an interface is by default PFC-enabled.
- PFC still allows you to configure lossless queues on a port to ensure no-drop handling of lossless traffic.

When you apply an input policy to an interface, an error message is displayed if:

- The PFC dot1p priorities result in more than two lossless port queues globally on the switch.
- You already enabled link-level flow control. PFC and link-level flow control cannot be enabled at the same time on an interface.

In a switch stack, configure all stacked ports with the same PFC configuration.

A DCB input policy for PFC applied to an interface may become invalid if you reconfigure the dot1p-queue mapping. This situation occurs when the new dot1p-queue assignment exceeds the maximum number (2) of lossless queues supported globally on the switch. In this case, all PFC configurations received from PFC-enabled peers are removed and resynchronized with the peer devices.

Traffic may be interrupted when you reconfigure PFC no-drop priorities in an input policy or reapply the policy to an interface.

Related Commands	dcb-input — creates a DCB input policy.

dcb-policy input stack-unit stack-ports all

Apply the specified DCB input policy on all ports of the switch stack or a single stacked switch.

Syntax	`dcb-policy input stack-unit {all | stack-unit-id} stack-ports all dcb-input-policy-name` To remove all DCB input policies applied to the stacked ports and rest the PFC to its default settings, use the `no dcb-policy input stack-unit all` command. To remove only the DCB input policies applied to the specified switch, use the `no dcb-policy input stack-unit` command.
Parameters	***stack-unit-id*** Enter the stack unit identification. ***dcb-input-policy-name*** Enter the policy name for the DCB input policy.
Defaults	none
Command Modes	CONFIGURATION
Command History	**Version 9.3.0.0** Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.
Usage Information	The `dcb-policy input stack-unit all` command overwrites any previous `dcb-policy input stack-unit stack-unit-id` configurations. Similarly, a `dcb-policy input stack-unit stack-unit-id` command overwrites any previous `dcb-policy input stack-unit all` configuration.
Related Commands	dcb-policy output stack-unit stack-ports all — applies the specified DCB output policy.

description

Enter a text description of the DCB policy (PFC input or ETS output).

Syntax	`description text` To remove the text description, use the `no description` command.
Parameters	***text*** Enter the description of the output policy. The maximum is 32 characters.
Defaults	none
Command Modes	• DCB INPUT POLICY • DCB OUTPUT POLICY

Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch

Related Commands

- dcb-input — creates a DCB PFC input policy.
- dcb-policy input — applies the output policy.
- dcb-output — creates a DCBETS output policy.
- dcb-policy output — applies the output policy.

pfc link-delay

Configure the link delay used to pause specified priority traffic.

Syntax `pfc link-delay value`

To remove the link delay, use the `no pfc link-delay` command.

Parameters

value	The range is (in quanta) from 712 to 65535. One quantum is equal to a 512-bit transmission.

Defaults **45556** quantum

Command Modes DCB INPUT POLICY

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information The minimum link delay must be greater than the round-trip transmission time a peer must honor a PFC pause frame multiplied by the number of PFC-enabled ingress ports.

Related Commands dcb-input — creates a DCB input policy.

pfc mode on

Enable the PFC configuration on the port so that the priorities are included in DCBX negotiation with peer PFC devices.

Syntax `pfc mode on`

To disable the PFC configuration, use the `no pfc mode on` command.

Defaults PFC mode is on.

Command Modes DCB INPUT POLICY

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

By applying a DCB input policy with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, also enable PFC on all DCB egress ports or configure the dot1p priority-queue assignment of PFC priorities to lossless queues (refer to `pfc no-drop queues`).

To disable PFC operation on an interface, enter the `no pfc mode on` command in DCB Input Policy Configuration mode. PFC is enabled and disabled as global DCB operation is enabled (`dcb-enable`) or disabled (`no dcb-enable`).

You cannot enable PFC and link-level flow control at the same time on an interface.

Related Commands

dcb-input — creates a DCB input policy.

pfc no-drop queues

Configure the port queues that still function as no-drop queues for lossless traffic.

Syntax

`pfc no-drop queues` *`queue-range`*

To remove the no-drop port queues, use the `no pfc no-drop queues` command.

Parameters

queue-range	Enter the queue range. Separate the queue values with a comma; specify a priority range with a dash; for example, `pfc no-drop queues 1,3` or `pfc no-drop queues 2-3`. The range is from 0 to 3.

Defaults

No lossless queues are configured.

Command Modes

INTERFACE

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

The maximum number of lossless queues globally supported on the switch is two.

- The following lists the dot1p priority-queue assignments.

dot1p Value in the Incoming Frame	Description heading
0	0
1	0
2	0
3	1
4	2
5	3

dot1p Value in the Incoming Frame	Description heading
6	3
7	3

pfc priority

Configure the CoS traffic to be stopped for the specified delay.

Syntax

`pfc priority` *`priority-range`*

To delete the pfc priority configuration, use the `no pfc priority` command.

Parameters

priority-range	Enter the 802.1p values of the frames to be paused. Separate the priority values with a comma; specify a priority range with a dash; for example, pfc priority 1,3,5-7. The range is from 0 to 7.

Defaults

none

Command Modes

DCB INPUT POLICY

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

You can enable any number of 802.1p priorities for PFC. Queues to which PFC priority traffic is mapped are lossless by default. Traffic may be interrupted due to an interface flap (going down and coming up) when you reconfigure the lossless queues for no-drop priorities in a PFC input policy and reapply the policy to an interface.

The maximum number of lossless queues supported on the switch is two.

A PFC peer must support the configured priority traffic (as DCBX detects) to apply PFC.

Related Commands

dcb-input — creates a DCB input policy.

show dcb

Displays the data center bridging status, the number of PFC-enabled ports, and the number of PFC-enabled queues.

Syntax

`show dcb [stack-unit` *`unit-number`*`]`

Parameters

unit number	Enter the DCB unit number. The range is from 0 to 5.

Command Modes	EXEC Privilege
Command History	**Version 9.3.0.0** Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.
Usage Information	Specify a stack-unit number on the Master switch in a stack.

Example

```
Dell# show dcb
stack-unit 0 port-set 0
      DCB Status : Enabled
    PFC Port Count : 56 (current), 56 (configured)
PFC Queue Count : 2 (current), 2 (configured)
```

show interface pfc

Displays the PFC configuration applied to ingress traffic on an interface, including priorities and link delay.

Syntax `show interface` *port-type slot/port* `pfc` {summary | detail}

Parameters

port-type slot/ port pfc	Enter the port-type slot and port PFC information.
{summary | detail}	Enter the keyword `summary` for a summary list of results or enter the keyword `detail` for a full list of results.

Command Modes	INTERFACE
Command History	**Version 9.3.0.0** Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Usage Information

To clear the PFC TLV counters, use the `clear pfc counters interface` *port-type slot/port* command.

The following describes the `show interface pfc summary` command shown in the following example.

Field	Description
Interface	Interface type with stack-unit and port number.
Admin mode is on Admin is enabled	PFC admin mode is on or off with a list of the configured PFC priorities. When the PFC admin mode is on, PFC advertisements are enabled to be sent and received from peers; received PFC configuration take effect. The admin operational status for a DCBX exchange of PFC configuration is enabled or disabled.
Remote is enabled, Priority list Remote	Operational status (enabled or disabled) of peer device for DCBX exchange of PFC configuration with a list of the configured PFC priorities. Willing status of peer device for

Field	Description
Willing Status is enabled	DCBX exchange (Willing bit received in PFC TLV): enabled or disable.
Local is enabled	DCBX operational status (enabled or disabled) with a list of the configured PFC priorities.
Operational status (local port)	Port state for current operational PFC configuration: • `Init`: Local PFC configuration parameters were exchanged with the peer. • `Recommend`: Remote PFC configuration parameters were received from the peer. • `Internally propagated`: PFC configuration parameters were received from the configuration source.
PFC DCBX Oper status	Operational status for the exchange of the PFC configuration on the local port: match (up) or mismatch (down).
State Machine Type	Type of state machine used for DCBX exchanges of the PFC parameters: Feature — for legacy DCBX versions; Symmetric — for an IEEE version.
TLV Tx Status	Status of the PFC TLV advertisements: enabled or disabled.
PFC Link Delay	Link delay (in quanta) used to pause specified priority traffic.
Application Priority TLV: FCOE TLV Tx Status	Status of FCoE advertisements in application priority TLVs from the local DCBX port: enabled or disabled.
Application Priority TLV: SCSI TLV Tx Status	Status of ISCSI advertisements in application priority TLVs from the local DCBX port: enabled or disabled.
Application Priority TLV: Local FCOE Priority Map	Priority bitmap the local DCBX port uses in FCoE advertisements in application priority TLVs.
Application Priority TLV: Local ISCSI Priority Map	Priority bitmap the local DCBX port uses in ISCSI advertisements in application priority TLVs.
Application Priority TLV: Remote FCOE Priority Map	Status of FCoE advertisements in application priority TLVs from the remote peer port: enabled or disabled.
Application Priority TLV: Remote ISCSI Priority Map	Status of iSCSI advertisements in application priority TLVs from the remote peer port: enabled or disabled.

Field	Description
PFC TLV Statistics: Input TLV pkts	Number of PFC TLVs received.
PFC TLV Statistics: Output TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error pkts	Number of PFC error packets received.
PFC TLV Statistics: Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: Pause Rx pkts	Number of PFC pause frames received.

Example (Summary)

```
Dell# show interfaces tengigabitethernet 0/49 pfc summary
Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled, Priority list is 4
  Remote Willing Status is enabled
  Local is enabled
  Oper status is Recommended
  PFC DCBX Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quantams
  Application Priority TLV Parameters :
  ------------------------------------
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8

Dell# show interfaces tengigabitethernet 0/49 pfc detail
  Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled
  Remote Willing Status is enabled
  Local is enabled
  Oper status is recommended
  PFC DCBX Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quanta
  Application Priority TLV Parameters :
  ------------------------------------
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8
  0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts,
  0 Pause Tx pkts, 0 Pause Rx pkts
```

show interface pfc statistics

Displays counters for the PFC frames received and transmitted (by dot1p priority class) on an interface.

Syntax

```
show interface port-type slot/port pfc statistics
```

Parameters

port-type	Enter the port type.
slot/port	Enter the slot/port number.

Command Modes

INTERFACE

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Example (Summary)

```
Dell#show interfaces te 0/3 pfc statistics
Interface TenGigabitEthernet 0/3

Priority Rx XOFF Frames Rx Total Frames Tx Total Frames
-----------------------------------------------------
0        0              0               0
1        0              0               0
2        0              0               0
3        0              0               0
4        0              0               0
5        0              0               0
6        0              0               0
7        0              0               0
```

show qos dcb-input

Displays the PFC configuration in a DCB input policy.

Syntax

```
show qos dcb-input [pfc-profile]
```

Parameters

pfc-profile	Enter the PFC profile.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Example

```
Dell(conf)# show qos dcb-input
dcb-input pfc-profile
  pfc link-delay 32
  pfc priority 0-1
dcb-input pfc-profile1
  no pfc mode on
  pfc priority 6-7
```

show stack-unit stack-ports pfc details

Displays the PFC configuration applied to ingress traffic on stacked ports, including PFC Operational mode on each unit with the configured priorities, link delay, and number of pause packets sent and received.

Syntax

```
show stack-unit {all | stack-unit} stack-ports {all | port-number} pfc details
```

Parameters

stack-unit	Enter the stack unit.
port-number	Enter the port number.

Command Modes

CONFIGURATION

Command History

Version 9.3.0.0	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch.

Example

```
Dell(conf)# show stack-unit all stack-ports all pfc details

stack unit 0 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts

stack unit 1 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

38

Data Center Bridging (DCB)

Data center bridging (DCB) is supported on the S4810 S4820T platform.

NOTE:
DCB is not supported when you use 10GBaseT ports for stacking.

Ethernet Enhancements in Data Center Bridging

The following section describes DCB.

In the Dell Networking operating system (FTOS) version 8.3.12.0, the S4810 S4820T system supports loading two *DCB_Config* files: *FCoE_DCB_Config* and *iSCSI_DCB_Config*.

These files are located in the root directory flash:/CONFIG_TEMPLATE. After copying the configuration files to the startup config and reloading the system.

The S4810 S4820T supports the following DCB features:

- Data center bridging exchange protocol (DCBx)
- Priority-based flow control (PFC)
- Enhanced transmission selection (ETS)

DCB refers to a set of IEEE Ethernet enhancements that provide data centers with a single, robust, converged network to support multiple traffic types, including local area network (LAN), server, and storage traffic. Through network consolidation, DCB results in reduced operational cost, simplified management, and easy scalability by avoiding the need to deploy separate application-specific networks.

For example, instead of deploying an Ethernet network for LAN traffic, include additional storage area networks (SANs) to ensure lossless Fibre Channel traffic, and a separate InfiniBand network for high-performance inter-processor computing within server clusters, only one DCB-enabled network is required in a data center. The Dell Networking switches that support a unified fabric and consolidate multiple network infrastructures use a single input/output (I/O) device called a converged network adapter (CNA).

A CNA is a computer input/output device that combines the functionality of a host bus adapter (HBA) with a network interface controller (NIC). Multiple adapters on different devices for several traffic types are no longer required.

Data center bridging satisfies the needs of the following types of data center traffic in a unified fabric:

LAN traffic	LAN traffic consists of many flows that are insensitive to latency requirements, while certain applications, such as streaming video, are more sensitive to latency. Ethernet functions as a best-effort network that may drop packets in the case of network congestion. IP networks rely on transport protocols (for example, TCP) for reliable data transmission with the associated cost of greater processing overhead and performance impact.
Storage traffic	Storage traffic based on Fibre Channel media uses the SCSI protocol for data transfer. This traffic typically consists of large data packets with a payload of 2K

	bytes that cannot recover from frame loss. To successfully transport storage traffic, data center Ethernet must provide no-drop service with lossless links.
InterProcess Communication (IPC) traffic	InterProcess Communication (IPC) traffic within high-performance computing clusters to share information. Server traffic is extremely sensitive to latency requirements.

To ensure lossless delivery and latency-sensitive scheduling of storage and service traffic and I/O convergence of LAN, storage, and server traffic over a unified fabric, IEEE data center bridging adds the following extensions to a classical Ethernet network:

- 802.1Qbb — Priority-based Flow Control (PFC)
- 802.1Qaz — Enhanced Transmission Selection (ETS)
- 802.1Qau — Congestion Notification
- Data Center Bridging Exchange (DCBx) protocol

NOTE: In FTOS version 8.3.12.0, only the PFC, ETS, and DCBx features are supported in data center bridging.

Priority-Based Flow Control

In a data center network, priority-based flow control (PFC) manages large bursts of one traffic type in multiprotocol links so that it does not affect other traffic types and no frames are lost due to congestion.

When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device. In this way, PFC ensures that PFC-enabled priority traffic is not dropped by the switch.

PFC enhances the existing 802.3x pause and 802.1p priority capabilities to enable flow control based on 802.1p priorities (classes of service). Instead of stopping all traffic on a link (as performed by the traditional Ethernet pause mechanism), PFC pauses traffic on a link according to the 802.1p priority set on a traffic type. You can create lossless flows for storage and server traffic while allowing for loss in case of LAN traffic congestion on the same physical interface.

The following illustration shows how PFC handles traffic congestion by pausing the transmission of incoming traffic with dot1p priority 3.

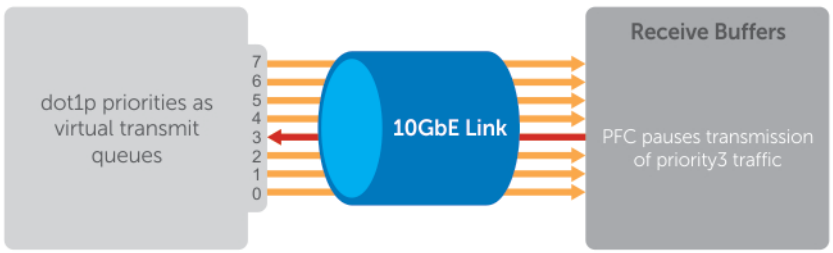

Figure 8. Priority-Based Flow Control

The system supports loading two DCB_Config files:

- FCoE_DCB_Config
- iSCSI_DCB_Config

In the FTOS operating system, PFC is implemented as follows:

- PFC supports buffering to receive data that continues to arrive on an interface while the remote system reacts to the PFC operation.
- PFC uses DCB MIB IEEE 802.1azd2.5 and PFC MIB IEEE 802.1bb-d2.2.
- PFC is supported on specified 802.1p priority traffic (dot1p 0 to 7) and is configured per interface. However, only two lossless queues are supported on an interface: one for Fibre Channel over Ethernet (FCoE) converged traffic and one for Internet Small Computer System Interface (iSCSI) storage traffic. Configure the same lossless queues on all ports.
- A dynamic threshold handles intermittent traffic bursts and varies based on the number of PFC priorities contending for buffers, while a static threshold places an upper limit on the transmit time of a queue after receiving a message to pause a specified priority. PFC traffic is paused only after surpassing both static and dynamic thresholds for the priority specified for the port.
- By default, PFC is enabled when you enable DCB. If you have not loaded FCoE_DCB_Config and iSCSI_DCB_Config, DCB is disabled. When you enable DCB globally, you cannot simultaneously enable link-level flow control.
- Buffer space is allocated and de-allocated only when you configure a PFC priority on the port.

Enhanced Transmission Selection

Enhanced transmission selection (ETS) supports optimized bandwidth allocation between traffic types in multiprotocol (Ethernet, FCoE, SCSI) links.

ETS allows you to divide traffic according to its 802.1p priority into different priority groups (traffic classes) and configure bandwidth allocation and queue scheduling for each group to ensure that each traffic type is correctly prioritized and receives its required bandwidth. For example, you can prioritize low-latency storage or server cluster traffic in a traffic class to receive more bandwidth and restrict best-effort LAN traffic assigned to a different traffic class.

Although you can configure strict-priority queue scheduling for a priority group, ETS introduces flexibility that allows the bandwidth allocated to each priority group to be dynamically managed according to the amount of LAN, storage, and server traffic in a flow. Unused bandwidth is dynamically allocated to prioritized priority groups. Traffic is queued according to its 802.1p priority assignment, while flexible bandwidth allocation and the configured queue-scheduling for a priority group is supported.

The following figure shows how ETS allows you to allocate bandwidth when different traffic types are classed according to 802.1p priority and mapped to priority groups.

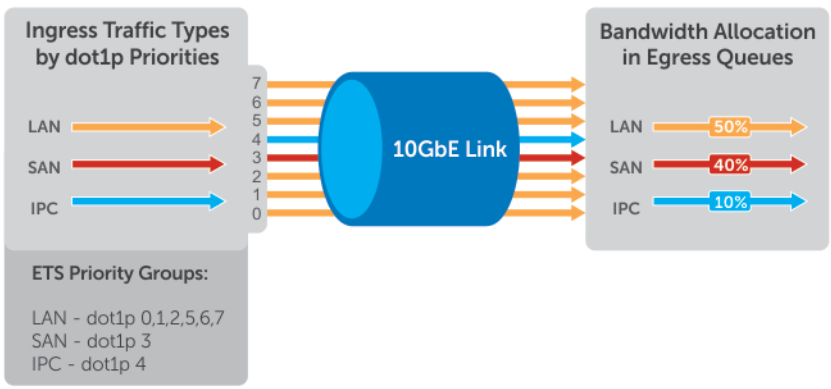

Figure 9. Enhanced Transmission Selection

The following table lists the traffic groupings ETS uses to select multiprotocol traffic for transmission.

Table 12. ETS Traffic Groupings

Traffic Groupings	Description
Priority group	A group of 802.1p priorities used for bandwidth allocation and queue scheduling. All 802.1p priority traffic in a group must have the same traffic handling requirements for latency and frame loss.
Group ID	A 4-bit identifier assigned to each priority group. The range is from 0 to 7.
Group bandwidth	Percentage of available bandwidth allocated to a priority group.
Group transmission selection algorithm (TSA)	Type of queue scheduling a priority group uses.

In FTOS, ETS is implemented as follows:

- ETS supports groups of 802.1p priorities that have:
  - PFC enabled or disabled
  - No bandwidth limit or no ETS processing
- Bandwidth allocated by the ETS algorithm is made available after strict-priority groups are serviced. Bandwidth is distributed in the following ways:
  - If bandwidth is not assigned to the priority groups, all available bandwidth is equally distributed among the priority groups. For example, if there are two priority groups, each will be assigned 50% of the available bandwidth; if there are four priority groups, each group is assigned 25% of the available bandwidth so that bandwidth use is always 100%.
  - If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups so that the sum of the bandwidth use is 100%.
  - If priority group bandwidth is less than 100%, all configured priority group bandwidth is incremented based on the configured percentage ratio until all priority group bandwidth use is 100%.
  - If priority group bandwidth use exceeds 100%, all configured priority group bandwidth is decremented based on the configured percentage ratio until all priority group bandwidth use is 100%.
  - If priority group bandwidth usage is greater than or equal to 100% and any default priority groups exist, then a minimum of 1% bandwidth use is assigned by decreasing 1% of bandwidth from the other priority groups until priority group bandwidth use is 100%.

  NOTE: You must configure at least one ETS priority group to a DCB output policy.
- For ETS traffic selection, an algorithm is applied to priority groups using:
  - Strict priority shaping
  - ETS shaping
  - (Credit-based shaping is not supported.)
- ETS uses the DCB MIB IEEE 802.1azd2.5.

Configuring DCB Maps and its Attributes

This topic contains the following sections that describe how to configure a DCB map, apply the configured DCB map to a port, configure PFC without a DCB map, and configure lossless queues. This functionality is supported S4810, S4820T, S6000, I/O Aggregator, and MXL platforms.

DCB Map: Configuration Procedure

A DCB map consists of PFC and ETS parameters. By default, PFC is not enabled on any 802.1p priority and ETS allocates equal bandwidth to each priority. To configure user-defined PFC and ETS settings, you must create a DCB map.

Step	Task	Command	Command Mode
1	Enter global configuration mode to create a DCB map or edit PFC and ETS settings.	dcb-map *name*	CONFIGURATION
2	Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group or whether priority group traffic should be handled with strict priority scheduling. You can enable PFC on a maximum of two priority queues on an interface. Enabling PFC for dot1p priorities makes the corresponding port queue lossless. The sum of all allocated bandwidth percentages in all groups in the DCB map must be 100%. Strict-priority traffic is serviced first. Afterwards, bandwidth allocated to other priority groups is made available and allocated according to the specified percentages. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups. Example: **priority-group 0 bandwidth 60 pfc off priority-group 1 bandwidth 20 pfc on priority-group 2 bandwidth 20 pfc on priority-group 4 strict-priority pfc off** Repeat this step to configure PFC and ETS traffic handling for each priority group.	priority-group *group_num* {**bandwidth** *percentage* | **strict-priority**} **pfc** {**on** | **off**}	DCB MAP
3	Specify the dot1p priority-to-priority group mapping for each priority. Priority-group range: 0 to 7. All priorities that map to the same queue must be in the same priority group. Leave a space between each priority group number. For example: **priority-pgid 0 0 0 1 2 4 4 4** in which priority group 0 maps to dot1p priorities 0, 1, and 2; priority group 1 maps to dot1p priority 3; priority group 2 maps to dot1p priority 4; priority group 4 maps to dot1p priorities 5, 6, and 7.	priority-pgid *dot1p0_group_num dot1p1_group_num dot1p2_group_num dot1p3_group_num dot1p4_group_num dot1p5_group_num dot1p6_group_num dot1p7_group_num*	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (**no priority pgid** command), the PFC and ETS parameters revert to their default values on the interfaces on which the

DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.

As a result, PFC and lossless port queues are disabled on 802.1p priorities, and all priorities are mapped to the same priority queue and equally share port bandwidth.

- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, first create a new DCB map with the desired PFC and ETS settings and apply the new map to the interfaces to override the previous DCB map settings. Then delete the original dot1p priority-priority group mapping.

  If you delete the dot1p priority-priority group mapping (**no priority pgid** command) before you apply the new DCB map, the default PFC and ETS parameters are applied on the interfaces. This change may create a DCB mismatch with peer DCB devices and interrupt network operation.

Applying a DCB Map on a Port

When you apply a DCB map with PFC enabled on an S6000 interface, a memory buffer for PFC-enabled priority traffic is automatically allocated. The buffer size is allocated according to the number of PFC-enabled priorities in the assigned map.

To apply a DCB map to an Ethernet port, follow these steps:

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	interface {**tengigabitEthernet** *slot*/*port* | **fortygigabitEthernet** *slot*/*port*}	CONFIGURATION
2	Apply the DCB map on the Ethernet port to configure it with the PFC and ETS settings in the map; for example: FTOS# **interface tengigabitEthernet 0/0** FTOS(config-if-te-0/0)# **dcb-map SAN_A_dcb_map1** Repeat Steps 1 and 2 to apply a DCB map to more than one port. You cannot apply a DCB map on an interface which has been already configured for PFC using the **pfc priority** command or which is already configured for lossless queues (**pfc no-drop queues** command).	dcb-map *name*	INTERFACE

Configuring PFC without a DCB Map

In a network topology that uses the default ETS bandwidth allocation (assigns equal bandwidth to each priority), you can also enable PFC for specified dot1p-priorities on individual interfaces without using a DCB map. This type of DCB configuration is useful on interfaces which require PFC for lossless traffic but do not transmit converged Ethernet traffic.

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	interface {tengigabitEthernet *slot*/***port*** | fortygigabitEthernet *slot*/***port***}	CONFIGURATION
2	Enable PFC on specified priorities. Range: 0-7. Default: None.	pfc priority *priority-range*	INTERFACE

Step	Task	Command	Command Mode
	Maximum number of loss less queues supported on an Ethernet port: 2. Separate priority values with a comma. Specify a priority range with a dash, for example: pfc priority 3,5-7 1. You cannot configure PFC using the **pfc priority** command on an interface on which a DCB map has been applied or which is already configured for lossless queues (**pfc no-drop queues** command).		

Configuring Lossless Queues

DCB also supports the manual configuration of lossless queues on an interface after you disable PFC mode in a DCB map and apply the map on the interface. The configuration of no-drop queues provides flexibility for ports on which PFC is not needed but lossless traffic should egress from the interface.

Lossless traffic egresses out the no-drop queues. Ingress 802.1p traffic from PFC-enabled peers is automatically mapped to the no-drop egress queues.

When configuring lossless queues on a port interface, take into account:

- By default, no lossless queues are configured on a port.
- A limit of two lossless queues are supported on a port. If the number of lossless queues configured exceeds the maximum supported limit per port (two), an error message displays. You must reconfigure the value to a smaller number of queues.
- If you configure lossless queues on an interface that already has a DCB map with PFC enabled (**pfc on**), an error message displays.

Step	Task	Command	Command Mode
1	Enter INTERFACE Configuration mode.	interface {**tengigabitEthernet** *slot*/*port* | **fortygigabitEthernet** *slot*/*port*}	CONFIGURATION
2	Open a DCB map and enter DCB map configuration mode.	dcb-map *name*	INTERFACE
3	Disable PFC.	no pfc mode on	DCB MAP
4	Return to interface configuration mode.	exit	DCB MAP
5	Apply the DCB map created to disable PFC operation on the interface	dcb-map {*name* | **default**}	INTERFACE
6	Configure the port queues that still function as no-drop queues for lossless traffic. For the dot1p-queue assignments, see Table 131. The maximum number of lossless queues globally supported on a port is 2. You cannot configure PFC no-drop queues on an interface on which a DCB map with PFC enabled has been applied or which is	**pfc no-drop queues** *queue-range*	INTERFACE

Step	Task	Command	Command Mode
	already configured for PFC using the **pfc priority** command. Range: 0-3. Separate queue values with a comma; specify a priority range with a dash; for example: pfc no-drop queues 1,3 or pfc no-drop queues 2-3 Default: No lossless queues are configured.		

Data Center Bridging: Default Configuration

This functionality is supported on the S6000 platform.

Before you configure PFC and ETS on an S5000 switch (see Configuring DCB Maps and its Attributes), take into account the following default settings:

DCB is enabled (see Enabling Data Center Bridging).

The PFC memory buffer supports up to 64 PFC-enabled ports and two lossless queues per port.

PFC and ETS are globally enabled by default:

The default dot1p priority-queue assignments are applied as follows:

802.1p value in incoming frame:	0	1	2	3	4	5	6	7
Egress queue assignment	0	0	0	1	2	3	3	3

PFC is not applied on specific dot1p priorities.

ETS: Equal bandwidth is assigned to each port queue and each dot1p priority in a priority group.

To configure PFC and ETS parameters on an S5000 interface, you must specify a PFC mode and ETS bandwidth allocation for a priority group and an 802.1p priority-to-priority group mapping in a DCB map (see Configuring PFC and ETS in a DCB Map). No default PFC and ETS settings are applied to Ethernet interfaces.

Configuring PFC and ETS in a DCB Map

An S6000 switch supports the use of a DCB map in which you configure priority-based flow control and enhanced transmission selection settings. To configure PFC and ETS parameters, you must apply a DCB map on an S6000 interface. This functionality is supported on the S6000 platform.

PFC Configuration Notes

Priority-based flow control (PFC) provides a flow control mechanism based on the 802.1p priorities in converged Ethernet traffic received on an interface and is enabled by default when you enable DCB. As an enhancement to the existing Ethernet pause mechanism, PFC stops traffic transmission for specified priorities (CoS values) without impacting other priority classes. Different traffic types are assigned to different priority classes.

When traffic congestion occurs, PFC sends a pause frame to a peer device with the CoS priority values of the traffic that needs to be stopped. DCBx provides the link-level exchange of PFC parameters between peer devices. PFC allows network administrators to create zero-loss links for SAN traffic that requires no-drop service, while at the same time retaining packet-drop congestion management for LAN traffic.

On an S6000 switch, PFC is enabled by default on Ethernet ports (**pfc mode on** command). You can configure PFC parameters using a DCB map or the **pfc priority** command in Interface configuration mode. For more information, see Configuring DCB Maps and its Attributes.

NOTE: DCB maps are supported only on physical Ethernet interfaces.

When you configure PFC in a DCB map:

- As soon as you apply a DCB map with PFC enabled on an interface, DCBx starts exchanging information with a peer. The IEEE802.1Qbb, CEE and CIN versions of PFC TLV are supported. DCBx also validates PFC configurations that are received in TLVs from peer devices.

By applying a DCB map with PFC enabled, you enable PFC operation on ingress port traffic. To achieve complete lossless handling of traffic, configure PFC priorities on all DCB egress ports.

- To remove a DCB map, including the PFC configuration it contains, use the **no dcb map** command in Interface configuration mode.
- To disable PFC operation on an interface, use the **no pfc mode on** command in DCB-Map configuration mode.
- Traffic may be interrupted when you reconfigure PFC no-drop priorities in a DCB map or re-apply the DCB map to an interface.
- For PFC to be applied, the configured priority traffic must be supported by a PFC peer (as detected by DCBx).
- If you apply a DCB map with PFC disabled (**pfc off**):
- You can enable link-level flow control on the interface (**flowcontrol rx on tx on** command; see Using Ethernet Pause Frames for Flow Control). To delete the DCB map, first disable link-level flow control. PFC is then automatically enabled on the interface because an interface is PFC-enabled by default.
- To ensure no-drop handling of lossless traffic, PFC allows you to configure lossless queues on a port (see Configuring DCB Maps and its Attributes).
- When you configure a DCB map, an error message displays if:
- The PFC dot1p priorities result in more than two lossless queues.
- When you apply a DCB map, an error message displays if:
- Link-level flow control is already enabled on an interface. You cannot enable PFC and link-level flow control at the same time on an interface.
- In a switch stack, configure all stacked ports with the same PFC configuration.
- FTOS allows you to change the default dot1p priority-queue assignments only if the change satisfies the following requirements in DCB maps already applied to S6000 interfaces:
- All 802.1p priorities mapped to the same queue must be in the same priority group.
- A maximum of two PFC-enabled, lossless queues are supported on an interface.

Otherwise the reconfiguration of a default dot1p-queue assignment is rejected.

- To ensure complete no-drop service, apply the same PFC parameters on all PFC-enabled peers.

PFC Prerequisites and Restrictions

On an S6000 switch, PFC is globally enabled by default, but not applied on specific 802.1p priorities. To enable PFC on 802.1p priorities, create a DCB map. For more information, see Configuring DCB Maps and its Attributes.

The following prerequisites and restrictions apply when you configure PFC in a DCB map:

- You can enable PFC on a maximum of two priority queues on an interface. Enabling PFC for dot1p priorities configures the corresponding port queue as lossless.

- You cannot enable PFC and link-level flow control at the same time on an interface.

ETS Configuration Notes

ETS provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic. Different traffic types have different service needs. Using ETS, you can create groups within an 802.1p priority class to configure different treatment for traffic with different bandwidth, latency, and best-effort needs.

When you configure ETS in a DCB map:

- The DCB map associates a priority group with a PFC operational mode (on or off) and an ETS scheduling and bandwidth allocation. You can apply a DCB map on multiple egress ports.
- Use the ETS configuration associated with 802.1p priority traffic in a DCB map in DCBx negotiation with ETS peers.
- Traffic in priority groups is assigned to strict-queue or weighted round-robin (WRR) scheduling in an ETS configuration and is managed using the ETS bandwidth-assignment algorithm. FTOS de-queues all frames of strict-priority traffic before servicing any other queues. A queue with strict-priority traffic can starve other queues in the same port.
- ETS-assigned bandwidth allocation and strict-priority scheduling apply only to data queues, not to control queues.
- FTOS supports hierarchical scheduling on an interface. FTOS control traffic is redirected to control queues as higher priority traffic with strict priority scheduling. After the control queues drain out, the remaining data traffic is scheduled to queues according to the bandwidth and scheduler configuration in the DCB map. The available bandwidth calculated by the ETS algorithm is equal to the link bandwidth after scheduling non-ETS higher-priority traffic.
- The configuration of bandwidth allocation and strict-queue scheduling is not supported at the same time for a priority group.
- **Bandwidth assignment**: By default, equal bandwidth is assigned to each dot1p priority in a priority group. To configure the bandwidth assigned to the port queues associated with dot1p priorities in a priority group, use the **bandwidth** *percentage* parameter. The sum of the bandwidth allocated to all priority groups in a DCB map must be 100% of the bandwidth on the link. You must allocate at least 1% of the total bandwidth to each priority group.
- **Scheduling of priority traffic**: dot1p priority traffic on the switch is scheduled to the current queue mapping. dot1p priorities within the same queue must have the same traffic properties and scheduling method.
- **ETS configuration error**: If an error occurs in an ETS configuration, the configuration is ignored and the scheduler and bandwidth allocation settings are reset to the ETS default value: 100% of available bandwidth is allocated to priority group 0 and bandwidth is equally assigned to each dot1p priority.

If an error occurs when a port receives a peer's ETS configuration, the port's configuration resets to the ETS configuration in the previously configured DCB map. If no DCB map was previously applied, the port resets to the default ETS parameters.

ETS Prerequisites and Restrictions

On an S6000 switch, ETS is enabled by default on Ethernet ports; equal bandwidth is assigned to each 802.1p priority. You can change the default ETS configuration only by using a DCB map. For more information, see Configuring DCB Maps and its Attributes.

The following prerequisites and restrictions apply when you configure ETS bandwidth allocation or strict-priority queuing in a DCB map:

- When allocating bandwidth or configuring strict-priority queuing for dot1p priorities in a priority group on a DCBx CIN interface, take into account the CIN bandwidth allocation (see Configuring Bandwidth Allocation for DCBx CIN) and dot1p-queue mapping.
- Although ETS bandwidth allocation or strict-priority queuing does not support weighted random early detection (WRED), explicit congestion notification (ECN), rate shaping, and rate limiting because these parameters are not negotiated by DCBx with peer devices, you can apply a QoS output policy with WRED and/or rate shaping on a DCBx CIN-enabled interface (see Configuring Port-based Rate Shaping and Weighted Random Early Detection). In this case, the WRED or rate shaping configuration in the QoS output policy must take into account the bandwidth allocation or queue scheduler configured in the DCB map.

Priority-Group Configuration Notes

When you configure priority groups in a DCB map:

- A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share the same latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.
- In a DCB map, each 802.1p priority must map to a priority group.
- The maximum number of priority groups supported in a DCB map on an interface is equal to the number of data queues (4) on the port. Each priority group can support more than one data queue.
- You can enable PFC on a maximum of two priority queues on an interface.
- If you configure more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

Applying a DCB Map in a Switch Stack

You must apply the same DCB map with PFC and ETS configuration to all stacked ports in a switch stack. You cannot apply different DCB maps to different stacked switches. This functionality is supported on the S6000 platform.

Entering the **no dcb-map stack-unit all** command removes all PFC and ETS settings applied to stacked ports from the DCB map and resets PFC and ETS to their default settings.

Task	Command	Command Mode
Apply the specified DCB map on all ports of the switch stack.	dcb-map stack-unit all stack-ports all *dcb-map-name*	CONFIGURATION

Data Center Bridging Exchange Protocol (DCBx)

The data center bridging exchange (DCBx) protocol is disabled by default on the S4810; ETS is also disabled.

DCBx allows a switch to automatically discover DCB-enabled peers and exchange configuration information. PFC and ETS use DCBx to exchange and negotiate parameters with peer devices. DCBx capabilities include:

- Discovery of DCB capabilities on peer-device connections.
- Determination of possible mismatch in DCB configuration on a peer link.

- Configuration of a peer device over a DCB link.

DCBx requires the link layer discovery protocol (LLDP) to provide the path to exchange DCB parameters with peer devices. Exchanged parameters are sent in organizationally specific TLVs in LLDP data units. For more information, refer to Link Layer Discovery Protocol (LLDP). The following LLDP TLVs are supported for DCB parameter exchange:

PFC parameters	PFC Configuration TLV and Application Priority Configuration TLV.
ETS parameters	ETS Configuration TLV and ETS Recommendation TLV.

Data Center Bridging in a Traffic Flow

The following figure shows how DCB handles a traffic flow on an interface.

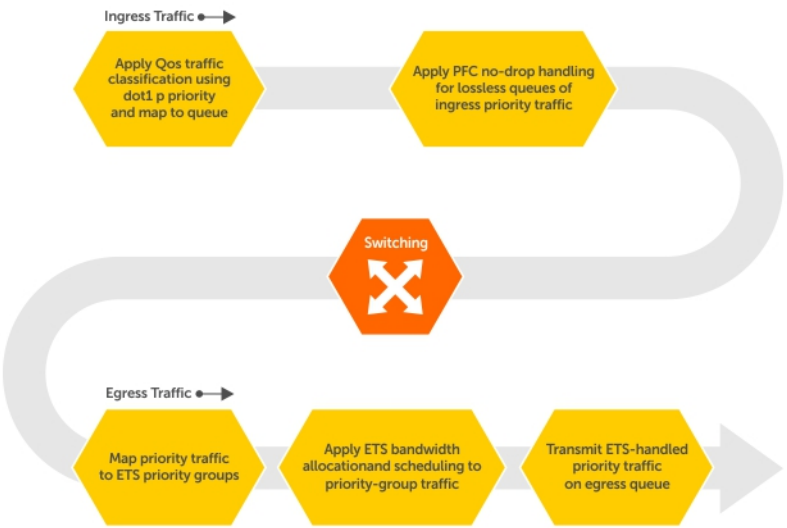
Ingress Traffic
Apply Qos traffic classification using dot1 p priority and map to queue
Apply PFC no-drop handling for lossless queues of ingress priority traffic
Switching
Egress Traffic
Map priority traffic to ETS priority groups
Apply ETS bandwidth allocationand scheduling to priority-group traffic
Transmit ETS-handled priority traffic on egress queue

Figure 10. DCB PFC and ETS Traffic Handling

Enabling Data Center Bridging

DCB is automatically configured when you configure FCoE or iSCSI optimization.

Data center bridging supports converged enhanced Ethernet (CEE) in a data center network. DCB is disabled by default. It must be enabled to support CEE.

- Priority-based flow control
- Enhanced transmission selection
- Data center bridging exchange protocol
- FCoE initialization protocol (FIP) snooping

DCB processes virtual local area network (VLAN)-tagged packets and dot1p priority values. Untagged packets are treated with a dot1p priority of 0.

For DCB to operate effectively, you can classify ingress traffic according to its dot1p priority so that it maps to different data queues. The dot1p-queue assignments used are shown in the following table.

To enable DCB, enable either the iSCSI optimization configuration or the FCoE configuration. For information to configure iSCSI optimization, refer to Enable and Disable iSCSI Optimization. For information to configure FCoE, refer to Step 1 in Configuring FIP Snooping.

To enable DCB with PFC buffers on a switch, enter the following commands, save the configuration, and reboot the system to allow the changes to take effect.

1. Enable DCB.
   CONFIGURATION mode
   ```
   dcb enable
   ```
2. Set PFC buffering on the DCB stack unit.
   CONFIGURATION mode
   ```
   dcb stack-unit all pfc-buffering pfc-ports 64 pfc-queues 2
   ```
   NOTE: To save the pfc buffering configuration changes, save the configuration and reboot the system.

NOTE: FTOS Behavior: DCB is not supported if you enable link-level flow control on one or more interfaces. For more information, refer to Ethernet Pause Frames.

QoS dot1p Traffic Classification and Queue Assignment

The following section describes QoS dot1P traffic classification and assignments.

DCB supports PFC, ETS, and DCBx to handle converged Ethernet traffic that is assigned to an egress queue according to the following QoS methods:

Honor dot1p	You can honor dot1p priorities in ingress traffic at the port or global switch level (refer to Default dot1p to Queue Mapping) using the `service-class dynamic dot1p` command in INTERFACE configuration mode.
Layer 2 class maps	You can use dot1p priorities to classify traffic in a class map and apply a service policy to an ingress port to map traffic to egress queues (refer to Policy-Based QoS Configurations).

NOTE: Dell Networking does not recommend mapping all ingress traffic to a single queue when using PFC and ETS. However, Dell Networking does recommend using Ingress traffic classification using the `service-class dynamic dot1p` command (honor dot1p) on all DCB-enabled interfaces. If you use L2 class maps to map dot1p priority traffic to egress queues, take into account the default dot1p-queue assignments in the following table and the maximum number of two lossless queues supported on a port (refer to Configuring Lossless Queues).

Although FTOS allows you to change the default dot1p priority-queue assignments (refer to Setting dot1p Priorities for Incoming Traffic), DCB policies applied to an interface may become invalid if you reconfigure dot1p-queue mapping. If the configured DCB policy remains valid, the change in the dot1p-queue assignment is allowed.

dot1p Value in the Incoming Frame	Egress Queue Assignment
0	0

dot1p Value in the Incoming Frame	Egress Queue Assignment
1	0
2	0
3	1
4	2
5	3
6	3
7	3

Configure Enhanced Transmission Selection

ETS provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic.

Different traffic types have different service needs. Using ETS, you can create groups within an 802.1p priority class to configure different treatment for traffic with different bandwidth, latency, and best-effort needs.

For example, storage traffic is sensitive to frame loss; interprocess communication (IPC) traffic is latency-sensitive. ETS allows different traffic types to coexist without interruption in the same converged link by:

- Allocating a guaranteed share of bandwidth to each priority group.
- Allowing each group to exceed its minimum guaranteed bandwidth if another group is not fully using its allotted bandwidth.

To configure ETS and apply an ETS output policy to an interface, you must:

1. Create a Quality of Service (QoS) output policy with ETS scheduling and bandwidth allocation settings.
2. Create a priority group of 802.1p traffic classes.
3. Configure a DCB output policy in which you associate a priority group with a QoS ETS output policy.
4. Apply the DCB output policy to an interface.

ETS Operation with DCBx

The following section describes DCBx negotiation with peer ETS devices.

In DCBx negotiation with peer ETS devices, ETS configuration is handled as follows:

- ETS TLVs are supported in DCBx versions CIN, CEE, and IEEE2.5.
- The DCBx port-role configurations determine the ETS operational parameters (refer to Configure a DCBx Operation).
- ETS configurations received from TLVs from a peer are validated.
- If there is a hardware limitation or TLV error:
  - DCBx operation on an ETS port goes down.
  - New ETS configurations are ignored and existing ETS configurations are reset to the previously configured ETS output policy on the port or to the default ETS settings if no ETS output policy was previously applied.

- ETS operates with legacy DCBx versions as follows:
  - In the CEE version, the priority group/traffic class group (TCG) ID 15 represents a non-ETS priority group. Any priority group configured with a scheduler type is treated as a strict-priority group and is given the priority-group (TCG) ID 15.
  - The CIN version supports two types of strict-priority scheduling:
    * Group strict priority: Use this to increase its bandwidth usage to the bandwidth total of the priority group and allow a single priority flow in a priority group. A single flow in a group can use all the bandwidth allocated to the group.
    * Link strict priority: Use this to increase to the maximum link bandwidth and allow a flow in any priority group.

  CIN supports only the dot1p priority-queue assignment in a priority group. To configure a dot1p priority flow in a priority group to operate with link strict priority, you configure: The dot1p priority for strict-priority scheduling (`strict-priority` command; Enabling Strict-Priority Queueing). The priority group for strict-priority scheduling (`scheduler strict` command; Creating a QoS ETS Output Policy).

  If you configure only the priority group in an ETS output policy or only the dot1p priority for strict-priority scheduling, the flow is handled with group strict priority.

Configuring Bandwidth Allocation for DCBx CIN

After you apply an ETS output policy to an interface, if the DCBx version used in your data center network is CIN, you may need to configure a QoS output policy to overwrite the default CIN bandwidth allocation.

This default setting divides the bandwidth allocated to each port queue equally between the dot1p priority traffic assigned to the queue.

For more information, refer to Allocating Bandwidth to Queue.

To create a QoS output policy that allocates different amounts of bandwidth to the different traffic types/dot1p priorities assigned to a queue and apply the output policy to the interface, follow these steps.

1. Create a QoS output policy.

   CONFIGURATION mode

   `qos-policy-output` *`output-policy-name`*

   The maximum 32 alphanumeric characters.
2. Configure the percentage of bandwidth to allocate to the dot1p priority/queue traffic in the associated L2 class map.

   QoS OUTPUT POLICY mode

   `bandwidth-percentage` *`percentage`*

   The default is **none**.
3. Repeat Step 2 to configure bandwidth percentages for other priority queues on the port.

   QoS OUTPUT POLICY mode

   `bandwidth-percentage` *`percentage`*
4. Exit QoS Output Policy Configuration mode.

   QoS OUTPUT POLICY mode

   `exit`
5. Enter INTERFACE Configuration mode.

   CONFIGURATION mode

   `interface type` *`slot/port`*

6. Apply the QoS output policy with the bandwidth percentage for specified priority queues to an egress interface.

   INTERFACE mode

   `service-policy output output-policy-name`

Configure a DCBx Operation

DCB devices use data center bridging exchange protocol (DCBx) to exchange configuration information with directly connected peers using the link layer discovery protocol (LLDP) protocol.

DCBx can detect the misconfiguration of a peer DCB device, and optionally, configure peer DCB devices with DCB feature settings to ensure consistent operation in a data center network.

DCBx is a prerequisite for using DCB features, such as priority-based flow control (PFC) and enhanced traffic selection (ETS), to exchange link-level configurations in a converged Ethernet environment. DCBx is also deployed in topologies that support lossless operation for FCoE or iSCSI traffic. In these scenarios, all network devices are DCBx-enabled (DCBx is enabled end-to-end). For more information about how these features are implemented and used, refer to:

- Configuring Priority-Based Flow Control
- Configure Enhanced Transmission Selection
- Configuring FIP Snooping

DCBx supports the following versions: CIN, CEE, and IEEE2.5.

Prerequisite: For DCBx, enable LLDP on all DCB devices.

DCBx Operation

DCBx performs the following operations:

- Discovers DCB configuration (such as PFC and ETS) in a peer device.
- Detects DCB mis-configuration in a peer device; that is, when DCB features are not compatibly configured on a peer device and the local switch. Mis-configuration detection is feature-specific because some DCB features support asymmetric configuration.
- Reconfigures a peer device with the DCB configuration from its configuration source if the peer device is willing to accept configuration.
- Accepts the DCB configuration from a peer if a DCBx port is in "willing" mode to accept a peer's DCB settings and then internally propagates the received DCB configuration to its peer ports.

DCBx Port Roles

To enable the auto-configuration of DCBx-enabled ports and propagate DCB configurations learned from peer DCBx devices internally to other switch ports, use the following DCBx port roles.

Auto-upstream — The port advertises its own configuration to DCBx peers and is *willing* to receive peer configuration. The port also propagates its configuration to other ports on the switch.

The first auto-upstream that is capable of receiving a peer configuration is elected as the configuration source. The elected configuration source then internally propagates the configuration to other auto-upstream and auto-downstream ports. A port that receives an internally propagated configuration overwrites its local configuration with the new parameter values. When an auto-upstream port

(besides the configuration source) receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBx peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.
  - The configuration received from a DCBx peer or from an internally propagated configuration is not stored in the switch's running configuration.
  - On a DCBx port in an auto-upstream role, the PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

Auto-downstream

The port advertises its own configuration to DCBx peers but is *not willing* to receive remote peer configuration. The port always accepts internally propagated configurations from a configuration source. An auto-downstream port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-downstream port receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBx peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.
  - The internally propagated configuration is not stored in the switch's running configuration.
  - On a DCBx port in an auto-downstream role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

Configuration source

The port is configured to serve as a source of configuration information on the switch. Peer DCB configurations received on the port are propagated to other DCBx auto-configured ports. If the peer configuration is compatible with a port configuration, DCBx is enabled on the port.

On a configuration-source port, the link with a DCBx peer is enabled when the port receives a DCB configuration that can be internally propagated to other auto-configured ports. The configuration received from a DCBx peer is not stored in the switch's running configuration. On a DCBx port that is the configuration source, all PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

Manual

The port is configured to operate only with administrator-configured settings and does not auto-configure with DCB settings received from a DCBx peer or from an internally propagated configuration from the configuration source. If you enable DCBx, ports in Manual mode advertise their configurations to peer devices but do not accept or propagate internal or external configurations. Unlike other user-

configured ports, the configuration of DCBx ports in Manual mode is saved in the running configuration.

On a DCBx port in a manual role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

When making a configuration change to a DCBx port in a Manual role, Dell Networking recommends shutting down the interface using the `shutdown` command, change the configuration, then re-activate the interface using the `no shutdown` command.

The default for the DCBx port role is **manual**.

NOTE: On a DCBx port, application priority TLV advertisements are handled as follows:

- The application priority TLV is transmitted only if the priorities in the advertisement match the configured PFC priorities on the port.
- On auto-upstream and auto-downstream ports:
  - If a configuration source is elected, the ports send an application priority TLV based on the application priority TLV received on the configuration-source port. When an application priority TLV is received on the configuration-source port, the auto-upstream and auto-downstream ports use the internally propagated PFC priorities to match against the received application priority. Otherwise, these ports use their locally configured PFC priorities in application priority TLVs.
  - If no configuration source is configured, auto-upstream and auto-downstream ports check to see that the locally configured PFC priorities match the priorities in a received application priority TLV.
- On manual ports, an application priority TLV is advertised only if the priorities in the TLV match the PFC priorities configured on the port.

DCB Configuration Exchange

The DCBx protocol supports the exchange and propagation of configuration information for the enhanced transmission selection (ETS) and priority-based flow control (PFC) DCB features.

DCBx uses the following methods to exchange DCB configuration parameters:

Asymmetric	DCB parameters are exchanged between a DCBx-enabled port and a peer port without requiring that a peer port and the local port use the same configured values for the configurations to be compatible. For example, ETS uses an asymmetric exchange of parameters between DCBx peers.
Symmetric	DCB parameters are exchanged between a DCBx-enabled port and a peer port but requires that each configured parameter value be the same for the configurations in order to be compatible. For example, PFC uses an symmetric exchange of parameters between DCBx peers.

Configuration Source Election

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port first checks to see if there is an active configuration source on the switch.

- If a configuration source already exists, the received peer configuration is checked against the local port configuration. If the received configuration is compatible, the DCBx marks the port as DCBx-enabled. If the configuration received from the peer is not compatible, a warning message is logged and the DCBx frame error counter is incremented. Although DCBx is operationally disabled, the port

keeps the peer link up and continues to exchange DCBx packets. If a compatible peer configuration is later received, DCBx is enabled on the port.

- If there is no configuration source, a port may elect itself as the configuration source. A port may become the configuration source if the following conditions exist:
  - No other port is the configuration source.
  - The port role is auto-upstream.
  - The port is enabled with link up and DCBx enabled.
  - The port has performed a DCBx exchange with a DCBx peer.
  - The switch is capable of supporting the received DCB configuration values through either a symmetric or asymmetric parameter exchange.

A newly elected configuration source propagates configuration changes received from a peer to the other auto-configuration ports. Ports receiving auto-configuration information from the configuration source ignore their current settings and use the configuration source information.

Propagation of DCB Information

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port acts as a DCBx client and checks if a DCBx configuration source exists on the switch.

- If a configuration source is found, the received configuration is checked against the currently configured values that are internally propagated by the configuration source. If the local configuration is compatible with the received configuration, the port is enabled for DCBx operation and synchronization.
- If the configuration received from the peer is not compatible with the internally propagated configuration used by the configuration source, the port is disabled as a client for DCBx operation and synchronization and a syslog error message is generated. The port keeps the peer link up and continues to exchange DCBx packets. If a compatible configuration is later received from the peer, the port is enabled for DCBx operation.

NOTE: DCB configurations internally propagated from a configuration source do not overwrite the configuration on a DCBx port in a manual role. When a configuration source is elected, all auto-upstream ports other than the configuration source are marked as *willing disabled*. The internally propagated DCB configuration is refreshed on all auto-configuration ports and each port may begin configuration negotiation with a DCBx peer again.

Auto-Detection and Manual Configuration of the DCBx Version

When operating in Auto-Detection mode (the `DCBx version auto` command), a DCBx port automatically detects the DCBx version on a peer port. Legacy CIN and CEE versions are supported in addition to the standard IEEE version 2.5 DCBx.

A DCBx port detects a peer version after receiving a valid frame for that version. The local DCBx port reconfigures to operate with the peer version and maintains the peer version on the link until one of the following conditions occurs:

- The switch reboots.
- The link is reset (goes down and up).
- User-configured CLI commands require the version negotiation to restart.
- The peer times out.
- Multiple peers are detected on the link.

If you configure a DCBx port to operate with a specific version (the `DCBx version {cee | cin | ieee-v2.5}` command in the Configuring DCBx), DCBx operations are performed according to the

configured version, including fast and slow transmit timers and message formats. If a DCBx frame with a different version is received, a syslog message is generated and the peer version is recorded in the peer status table. If the frame cannot be processed, it is discarded and the discard counter is incremented.

NOTE: Because DCBx TLV processing is best effort, it is possible that CIN frames may be processed when DCBx is configured to operate in CEE mode and vice versa. In this case, the unrecognized TLVs cause the unrecognized TLV counter to increment, but the frame is processed and is not discarded.

Legacy DCBx (CIN and CEE) supports the DCBx control state machine that is defined to maintain the sequence number and acknowledge the number sent in the DCBx control TLVs.

DCBx Example

The following figure shows how to use DCBx.

The external 40GbE ports on the base module (ports 33 and 37) of two switches are used for uplinks configured as DCBx auto-upstream ports. The S4810 S4820T is connected to third-party, top-of-rack (ToR) switches through 40GbE uplinks. The ToR switches are part of a Fibre Channel storage network.

The internal ports (ports 1-32) connected to the 10GbE backplane are configured as auto-downstream ports.

On the S4810 S4820T, PFC and ETS use DCBx to exchange link-level configuration with DCBx peer devices.

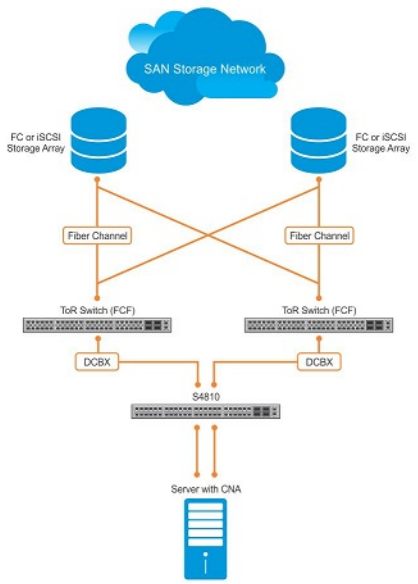

Figure 11. DCBx Sample Topology

DCBx Prerequisites and Restrictions

The following prerequisites and restrictions apply when you configure DCBx operation on a port:

- For DCBx, on a port interface, enable LLDP in both Send (TX) and Receive (RX) mode (the `protocol lldp mode` command; refer to the example in CONFIGURATION versus INTERFACE Configurations in the Link Layer Discovery Protocol (LLDP) chapter). If multiple DCBx peer ports are detected on a local DCBx interface, LLDP is shut down.

- The CIN version of DCBx supports only PFC, ETS, and FCOE; it does not support iSCSI, backward congestion management (BCN), logical link down (LLDF), and network interface virtualization (NIV).

Configuring DCBx

To configure DCBx, follow these steps.

For DCBx, to advertise DCBx TLVs to peers, enable LLDP. For more information, refer to Link Layer Discovery Protocol (LLDP).

Configure DCBx operation at the interface level on a switch or globally on the switch. To configure the S4810 S4820T system for DCBx operation in a data center network, you must:

1. Configure ToR- and FCF-facing interfaces as auto-upstream ports.
2. Configure server-facing interfaces as auto-downstream ports.
3. Configure a port to operate in a configuration-source role.
4. Configure ports to operate in a manual role.

1. Enter INTERFACE Configuration mode.

CONFIGURATION mode

```
interface type slot/port
```

2. Enter LLDP Configuration mode to enable DCBx operation.

INTERFACE mode

```
[no] protocol lldp
```

3. Configure the DCBx version used on the interface, where: `auto` configures the port to operate using the DCBx version received from a peer.

PROTOCOL LLDP mode

```
[no] DCBx version {auto | cee | cin | ieee-v2.5}
```

- `cee`: configures the port to use CEE (Intel 1.01).
- `cin`: configures the port to use Cisco-Intel-Nuova (DCBx 1.0).
- `ieee-v2.5`: configures the port to use IEEE 802.1Qaz (Draft 2.5).

The default is **Auto**.

4. Configure the DCBx port role the interface uses to exchange DCB information.

PROTOCOL LLDP mode

```
[no] DCBx port-role {config-source | auto-downstream | auto-upstream |
manual}
```

- `auto-upstream`: configures the port to receive a peer configuration. The configuration source is elected from auto-upstream ports.
- `auto-downstream`: configures the port to accept the internally propagated DCB configuration from a configuration source.
- `config-source`: configures the port to serve as the configuration source on the switch.
- `manual`: configures the port to operate only on administer-configured DCB parameters. The port does not accept a DCB configuration received from a peer or a local configuration source.

The default is **Manual**.

5. **On manual ports only**: Configure the PFC and ETS TLVs advertised to DCBx peers.

   PROTOCOL LLDP mode

   ```
   [no] advertise DCBx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]
   ```

   - `ets-conf`: enables the advertisement of ETS Configuration TLVs.
   - `ets-reco`: enables the advertisement of ETS Recommend TLVs.
   - `pfc enables`: the advertisement of PFC TLVs.

   The default is All PFC and ETS TLVs are advertised.

   NOTE: You can configure the transmission of more than one TLV type at a time; for example, advertise `DCBx-tlv ets-conf ets-reco`. You can enable ETS recommend TLVs (`ets-reco`) only if you enable ETS configuration TLVs (`ets-conf`).

   To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-tlv pfc ets-reco`.

6. **On manual ports only**: Configure the Application Priority TLVs advertised on the interface to DCBx peers.

   PROTOCOL LLDP mode

   ```
   [no] advertise DCBx-appln-tlv {fcoe | iscsi}
   ```

   - `fcoe`: enables the advertisement of FCoE in Application Priority TLVs.
   - `iscsi`: enables the advertisement of iSCSI in Application Priority TLVs.

   The default is Application Priority TLVs are enabled to advertise FCoE and iSCSI.

   NOTE: To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-appln-tlv iscsi`.

   For information about how to use FCoE and iSCSI, refer to FCoE Initialization Protocol (FIP) Snooping and iSCSI Optimization.

To verify the DCBx configuration on a port, use the `show interface DCBx detail` command.

Configuring DCBx Globally on the Switch

To globally configure the DCBx operation on a switch, follow these steps.

1. Enter Global Configuration mode.

   EXEC PRIVILEGE mode

   ```
   configure
   ```

2. Enter LLDP Configuration mode to enable DCBx operation.

   CONFIGURATION mode

   ```
   [no] protocol lldp
   ```

3. Configure the DCBx version used on all interfaces not already configured to exchange DCB information.

   PROTOCL LLDP mode

   `[no] DCBx version {auto | cee | cin | ieee-v2.5}`

   - `auto`: configures all ports to operate using the DCBx version received from a peer.
   - `cee`: configures a port to use CEE (Intel 1.01). cin configures a port to use Cisco-Intel-Nuova (DCBx 1.0).
   - `ieee-v2.5`: configures a port to use IEEE 802.1Qaz (Draft 2.5).

   The default is **Auto**.

   NOTE: To configure the DCBx port role the interfaces use to exchange DCB information, use the `DCBx port-role` command in INTERFACE Configuration mode (Step 3).

4. Configure the PFC and ETS TLVs that advertise on unconfigured interfaces with a manual port-role.

   PROTOCOL LLDP mode

   `[no] advertise DCBx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]`

   - `ets-conf`: enables transmission of ETS Configuration TLVs.
   - `ets-reco`: enables transmission of ETS Recommend TLVs.
   - `pfc`: enables transmission of PFC TLVs.

   NOTE: You can configure the transmission of more than one TLV type at a time. You can only enable ETS recommend TLVs (`ets-reco`) if you enable ETS configuration TLVs (`ets-conf`). To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-tlv pfc ets-reco`.

   The default is All TLV types are enabled.

5. Configure the Application Priority TLVs that advertise on unconfigured interfaces with a manual port-role.

   PROTOCOL LLDP mode

   `[no] advertise DCBx-appln-tlv {fcoe | iscsi}`

   - `fcoe`: enables the advertisement of FCoE in Application Priority TLVs.
   - `iscsi`: enables the advertisement of iSCSI in Application Priority TLVs.

   The default is Application Priority TLVs are enabled and advertise FCoE and iSCSI.

   NOTE: To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-appln-tlv iscsi`.

   For information about how to use FCoE and iSCSI, refer to FCoE Initialization Protocol (FIP) Snooping and iSCSI Optimization.

6. Configure the FCoE priority advertised for the FCoE protocol in Application Priority TLVs.

   PROTOCOL LLDP mode

   `[no] fcoe priority-bits` *`priority-bitmap`*

   The priority-bitmap range is from 1 to FF.

   The default is **0x8**.

7. Configure the iSCSI priority advertised for the iSCSI protocol in Application Priority TLVs.

   PROTOCOL LLDP mode

   `[no] iscsi priority-bits priority-bitmap`

   The priority-bitmap range is from 1 to FF.

   The default is **0x10**.

DCBx Error Messages

The following syslog messages appear when an error in DCBx operation occurs.

```
LLDP_MULTIPLE_PEER_DETECTED: DCBx is operationally disabled after detecting
more than one DCBx
peer on the port interface.
```

```
LLDP_PEER_AGE_OUT: DCBx is disabled as a result of LLDP timing out on a DCBx
peer interface.
```

```
DSM_DCBx_PEER_VERSION_CONFLICT: A local port expected to receive the IEEE, CIN,
or CEE version
in a DCBx TLV from a remote peer but received a different, conflicting DCBx
version.
```

```
DSM_DCBx_PFC_PARAMETERS_MATCH and DSM_DCBx_PFC_PARAMETERS_MISMATCH: A local
DCBx port received
a compatible (match) or incompatible (mismatch) PFC configuration from a peer.
```

```
DSM_DCBx_ETS_PARAMETERS_MATCH and DSM_DCBx_ETS_PARAMETERS_MISMATCH: A local
DCBx port received
a compatible (match) or incompatible (mismatch) ETS configuration from a peer.
```

```
LLDP_UNRECOGNISED_DCBx_TLV_RECEIVED: A local DCBx port received an unrecognized
DCBx TLV from
a peer.
```

Debugging DCBx on an Interface

To enable DCBx debug traces for all or a specific control paths, use the following command.

- Enable DCBx debugging.

  EXEC PRIVILEGE mode

  `debug DCBx {all | auto-detect-timer | config-exchng | fail | mgmt | resource | sem | tlv}`

  - `all`: enables all DCBx debugging operations.
  - `auto-detect-timer`: enables traces for DCBx auto-detect timers.
  - `config-exchng`: enables traces for DCBx configuration exchanges.
  - `fail`: enables traces for DCBx failures.
  - `mgmt`: enables traces for DCBx management frames.
  - `resource`: enables traces for DCBx system resource frames.
  - `sem`: enables traces for the DCBx state machine.
  - `tlv`: enables traces for DCBx TLVs.

Verifying the DCB Configuration

To display DCB configurations, use the following `show` commands.

Table 13. Displaying DCB Configurations

Command	Output
`show dot1p-queue mapping`	Displays the current 802.1p priority-queue mapping.
`show dcb [stack-unit` *`unit-number`*`]`	Displays the data center bridging status, number of PFC-enabled ports, and number of PFC-enabled queues. On the master switch in a stack, you can specify a stack-unit number. The range is from 0 to 5.
`show qos dcb-input [pfc-profile]`	Displays the PFC configuration in a DCB input policy.
`show qos dcb-output [ets-profile]`	Displays the ETS configuration in a DCB output policy.
`show qos priority-groups`	Displays the ETS priority groups configured on the switch, including the 802.1p priority classes and ID of each group.
`show interface` *`port-type slot/port`* `pfc {summary | detail}`	Displays the PFC configuration applied to ingress traffic on an interface, including priorities and link delay. To clear PFC TLV counters, use the `clear pfc counters interface` *`port-type slot/port`* command.
`show interface` *`port-type slot/port`* `pfc statistics`	Displays counters for the PFC frames received and transmitted (by dot1p priority class) on an interface.
`show interface` *`port-type slot/port`* `ets {summary | detail}`	Displays the ETS configuration applied to egress traffic on an interface, including priority groups with priorities and bandwidth allocation. To clear ETS TLV counters, enter the `clear ets counters interface` *`port-type slot/port`* command.
`show interface` *`port-type slot/port`* `DCBx detail`	Plays the DCBx configuration on an interface.
`show stack-unit {`*`0-11`*` | all} stack ports all pfc details`	Displays the PFC configuration applied to ingress traffic on stack-links, including priorities and link delay.
`show stack-unit {`*`0-11`*` | all} stack ports all ets details`	Displays the ETS configuration applied to ingress traffic on stack-links, including priorities and link delay.

Example of the `show dot1p-queue mapping` Command

```
FTOS(conf)# show dot1p-queue-mapping
Dot1p Priority: 0 1 2 3 4 5 6 7
Queue         : 0 0 0 1 2 3 3 3
```

Example of the `show dcb` Command

```
FTOS# show dcb
stack-unit 0 port-set 0
     DCB Status : Enabled
 PFC Port Count : 56 (current), 56 (configured)
PFC Queue Count : 2 (current),  2 (configured)
```

Example of the `show qos dcb-input` Command

```
FTOS(conf)# show qos dcb-input
dcb-input pfc-profile
  pfc link-delay 32
  pfc priority 0-1
dcb-input pfc-profile1
  no pfc mode on
  pfc priority 6-7
```

Example of the `show qos dcb-output` Command

```
FTOS# show qos dcb-output
dcb-output ets
  priority-group san qos-policy san
  priority-group ipc qos-policy ipc
  priority-group lan qos-policy lan
```

Example of the `show qos priority-groups` Command

```
FTOS#show qos priority-groups
priority-group ipc
  priority-list 4
  set-pgid 2
```

Example of the `show interfaces pfc summary` Command

```
FTOS# show interfaces tengigabitethernet 0/49 pfc summary
Interface TenGigabitEthernet 0/49
    Admin mode is on
    Admin is enabled
    Remote is enabled, Priority list is 4
    Remote Willing Status is enabled
    Local is enabled
    Oper status is Recommended
    PFC DCBx Oper status is Up
    State Machine Type is Feature
    TLV Tx Status is enabled
    PFC Link Delay 45556 pause quantams
    Application Priority TLV Parameters :
    ---------------------------------------
    FCOE TLV Tx Status is disabled
    ISCSI TLV Tx Status is disabled
    Local FCOE PriorityMap is 0x8
    Local ISCSI PriorityMap is 0x10
    Remote FCOE PriorityMap is 0x8
    Remote ISCSI PriorityMap is 0x8

FTOS# show interfaces tengigabitethernet 0/49 pfc detail
Interface TenGigabitEthernet 0/49
    Admin mode is on
```

```
    Admin is enabled
    Remote is enabled
    Remote Willing Status is enabled
    Local is enabled
    Oper status is recommended
    PFC DCBx Oper status is Up
    State Machine Type is Feature
    TLV Tx Status is enabled
    PFC Link Delay 45556 pause quanta
    Application Priority TLV Parameters :
    ------------------------------------
    FCOE TLV Tx Status is disabled
    ISCSI TLV Tx Status is disabled
    Local FCOE PriorityMap is 0x8
    Local ISCSI PriorityMap is 0x10
    Remote FCOE PriorityMap is 0x8
    Remote ISCSI PriorityMap is 0x8

0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts, 0 Pause Tx pkts, 0 Pause Rx
pkts
```

The following table describes the `show interface pfc summary` command fields.

Table 14. `show interface pfc summary` Command Description

Fields	Description
Interface	Interface type with stack-unit and port number.
Admin mode is on; Admin is enabled	PFC Admin mode is on or off with a list of the configured PFC priorities . When PFC admin mode is on, PFC advertisements are enabled to be sent and received from peers; received PFC configuration takes effect. The admin operational status for a DCBx exchange of PFC configuration is enabled or disabled.
Remote is enabled; Priority list Remote Willing Status is enabled	Operational status (enabled or disabled) of peer device for DCBx exchange of PFC configuration with a list of the configured PFC priorities. Willing status of peer device for DCBx exchange (Willing bit received in PFC TLV): enabled or disabled.
Local is enabled	DCBx operational status (enabled or disabled) with a list of the configured PFC priorities
Operational status (local port)	DCBx operational status (enabled or disabled) with a list of the configured PFC priorities. Port state for current operational PFC configuration: • Init: Local PFC configuration parameters were exchanged with peer. • Recommend: Remote PFC configuration parameters were received from peer. • Internally propagated: PFC configuration parameters were received from configuration source.

Fields	Description
PFC DCBx Oper status	Operational status for exchange of PFC configuration on local port: match (up) or mismatch (down).
State Machine Type	Type of state machine used for DCBx exchanges of PFC parameters: • Feature: for legacy DCBx versions • Symmetric: for an IEEE version
TLV Tx Status	Status of PFC TLV advertisements: enabled or disabled.
PFC Link Delay	Link delay (in quanta) used to pause specified priority traffic.
Application Priority TLV: FCOE TLV Tx Status	Status of FCoE advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: ISCSI TLV Tx Status	Status of ISCSI advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: Local FCOE Priority Map	Priority bitmap used by local DCBx port in FCoE advertisements in application priority TLVs.
Application Priority TLV: Local ISCSI Priority Map	Priority bitmap used by local DCBx port in ISCSI advertisements in application priority TLVs.
Application Priority TLV: Remote FCOE Priority Map	Status of FCoE advertisements in application priority TLVs from remote peer port: enabled or disabled.
Application Priority TLV: Remote ISCSI Priority Map	Status of iSCSI advertisements in application priority TLVs from remote peer port: enabled or disabled.
PFC TLV Statistics: Input TLV pkts	Number of PFC TLVs received.
PFC TLV Statistics: Output TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error pkts	Number of PFC error packets received.
PFC TLV Statistics: Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: Pause Rx pkts	Number of PFC pause frames received

Example of the `show interface pfc statistics` Command

```
FTOS#show interfaces te 0/0 pfc statistics
Interface TenGigabitEthernet 0/0
Priority Received PFC Frames Transmitted PFC Frames
-------- ------------------- ----------------------
0        0                   0
1        0                   0
2        0                   0
3        0                   0
4        0                   0
```

```
5        0                   0
6        0                   0
7        0                   0
```

Example of the `show interface ets summary` Command

```
FTOS(conf-qos-policy-out-ets)#do sho int te 0/3 ets su

Interface TenGigabitEthernet 0/3
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters :
------------------
Admin is enabled

TC-grp   Priority#   Bandwidth   TSA
-----------------------------------------------
0                    -           -
1        0,1,2       100%        ETS
2        3           0 %         SP
3        4,5,6,7     0 %         SP
4                    -           -
5                    -           -
6                    -           -
7                    -           -

Remote Parameters :
-------------------
Remote is disabled

Local Parameters :
------------------
Local is enabled

TC-grp    Priority#     Bandwidth    TSA
-----------------------------------------------
0                       -            -
1         0,1,2         100%         ETS
2         3             0 %          SP
3         4,5,6,7       0 %          SP
4                       -            -
5                       -            -
6                       -            -
7                       -            -

Oper status is init
ETS DCBx Oper status is Down
State Machine Type is Asymmetric
Conf TLV Tx Status is enabled
Reco TLV Tx Status is enabled

0 Input Conf TLV Pkts, 1955 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
0 Input Reco TLV Pkts, 1955 Output Reco TLV Pkts, 0 Error Reco TLV Pkts
FTOS(conf)# show interfaces tengigabitethernet 0/0 ets detail
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :
------------------
Admin is enabled
TC-grp Priority# Bandwidth TSA
```

```
0 0,1,2,3,4,5,6,7 100% ETS
1 0% ETS
2 0% ETS
3 0% ETS
4 0% ETS
5 0% ETS
6 0% ETS
7 0% ETS
Priority# Bandwidth TSA
0 13% ETS
1 13% ETS
2 13% ETS
3 13% ETS
4 12% ETS
5 12% ETS
6 12% ETS
7 12% ETS
Remote Parameters:
-------------------
Remote is disabled
Local Parameters :
------------------
Local is enabled
TC-grp Priority# Bandwidth TSA
0 0,1,2,3,4,5,6,7 100% ETS
1 0% ETS
2 0% ETS
3 0% ETS
4 0% ETS
5 0% ETS
6 0% ETS
7 0% ETS
Priority# Bandwidth TSA
0 13% ETS
1 13% ETS
2 13% ETS
3 13% ETS
4 12% ETS
5 12% ETS
6 12% ETS
7 12% ETS
Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
0T LIVnput Traffic Class TLV Pkts, 0 Output Traffic Class TLV Pkts, 0 Error
Traffic Class
Pkts
```

Example of the `show interface ets detail` Command

```
FTOS(conf)# show interfaces tengigabitethernet 0/0 ets detail
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :
------------------
Admin is enabled
TC-grp    Priority#         Bandwidth    TSA
0         0,1,2,3,4,5,6,7   100%         ETS
1                           0%           ETS
2                           0%           ETS
3                           0%           ETS
```

```
4                             0%            ETS
5                             0%            ETS
6                             0%            ETS
7                             0%            ETS

Priority# Bandwidth TSA
0                             13%           ETS
1                             13%           ETS
2                             13%           ETS
3                             13%           ETS
4                             12%           ETS
5                             12%           ETS
6                             12%           ETS
7                             12%           ETS
Remote Parameters:
-------------------
Remote is disabled

Local Parameters :
------------------
Local is enabled
TC-grp    Priority#           Bandwidth     TSA
0         0,1,2,3,4,5,6,7     100%          ETS
1                             0%            ETS
2                             0%            ETS
3                             0%            ETS
4                             0%            ETS
5                             0%            ETS
6                             0%            ETS
7                             0%            ETS

Priority#                     Bandwidth     TSA
0                             13%           ETS
1                             13%           ETS
2                             13%           ETS
3                             13%           ETS
4                             12%           ETS
5                             12%           ETS
6                             12%           ETS
7                             12%           ETS
Oper status is init
Conf TLV Tx Status is disabled
Traffic Class TLV Tx Status is disabled
0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts
0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV Pkts, 0 Error
Traffic Class TLV
Pkts
```

The following table describes the `show interface ets detail` command fields.

Table 15. `show interface ets detail` Command Description

Field	Description
Interface	Interface type with stack-unit and port number.
Max Supported TC Group	Maximum number of priority groups supported.
Number of Traffic Classes	Number of 802.1p priorities currently configured.
Admin mode	ETS mode: on or off. When on, the scheduling and bandwidth allocation configured in an ETS output policy or received in a

Field	Description
	DCBx TLV from a peer can take effect on an interface.
Admin Parameters	ETS configuration on local port, including priority groups, assigned dot1p priorities, and bandwidth allocation.
Remote Parameters	ETS configuration on remote peer port, including Admin mode (enabled if a valid TLV was received or disabled), priority groups, assigned dot1p priorities, and bandwidth allocation. If the ETS Admin mode is enabled on the remote port for DCBx exchange, the Willing bit received in ETS TLVs from the remote peer is included.
Local Parameters	ETS configuration on local port, including Admin mode (enabled when a valid TLV is received from a peer), priority groups, assigned dot1p priorities, and bandwidth allocation.
Operational status (local port)	Port state for current operational ETS configuration: • Init: Local ETS configuration parameters were exchanged with peer. • Recommend: Remote ETS configuration parameters were received from peer. • Internally propagated: ETS configuration parameters were received from configuration source.
ETS DCBx Oper status	Operational status of ETS configuration on local port: match or mismatch.
State Machine Type	Type of state machine used for DCBx exchanges of ETS parameters: • Feature: for legacy DCBx versions • Asymmetric: for an IEEE version
Conf TLV Tx Status	Status of ETS Configuration TLV advertisements: enabled or disabled.
ETS TLV Statistic: Input Conf TLV pkts	Number of ETS Configuration TLVs received.
ETS TLV Statistic: Output Conf TLV pkts	Number of ETS Configuration TLVs transmitted.
ETS TLV Statistic: Error Conf TLV pkts	Number of ETS Error Configuration TLVs received.

Example of the `show stack-unit all stack-ports all pfc details` Command

```
FTOS(conf)# show stack-unit all stack-ports all pfc details

stack unit 0 stack-port all
    Admin mode is On
    Admin is enabled, Priority list is 4-5
    Local is enabled, Priority list is 4-5
```

```
    Link Delay 45556 pause quantum
    0 Pause Tx pkts, 0 Pause Rx pkts

stack unit 1 stack-port all
    Admin mode is On
    Admin is enabled, Priority list is 4-5
    Local is enabled, Priority list is 4-5
    Link Delay 45556 pause quantum
    0 Pause Tx pkts, 0 Pause Rx pkts
```

Example of the `show stack-unit all stack-ports all ets details` Command

```
FTOS(conf)# show stack-unit all stack-ports all ets details
Stack unit 0 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1

Admin mode is on
Admin Parameters:
--------------------
Admin is enabled
TC-grp    Priority#        Bandwidth    TSA
------------------------------------------------
0         0,1,2,3,4,5,6,7  100%         ETS
1                          -            -
2                          -            -
3                          -            -
4                          -            -
5                          -            -
6                          -            -
7                          -            -
8                          -            -

Stack unit 1 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
Admin mode is on
Admin Parameters:
--------------------
Admin is enabled
TC-grp    Priority#        Bandwidth    TSA
------------------------------------------------
0         0,1,2,3,4,5,6,7  100%         ETS
1                          -            -
2                          -            -
3                          -            -
4                          -            -
5                          -            -
6                          -            -
7                          -            -
8                          -            -
```

Example of the `show interface DCBx detail` Command (IEEE)

```
FTOS(conf-if-te-0/17-lldp)#do sho int te 2/12 dc d

E-ETS Configuration TLV enabled             e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled            r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled             p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled     f-Application Priority for FCOE
disabled
I-Application priority for iSCSI enabled    i-Application Priority for iSCSI
disabled
-------------------------------------------------------------------------------
```

```
----------

Interface TenGigabitEthernet 2/12
  Remote Mac Address 00:01:e8:8a:df:a0
    Port Role is Manual
  DCBx Operational Status is Enabled
  Is Configuration Source? FALSE
  Local DCBx Compatibility mode is IEEEv2.5
  Local DCBx Configured mode is IEEEv2.5
  Peer Operating version is IEEEv2.5
  Local DCBx TLVs Transmitted: ERPFi
  1 Input PFC TLV pkts, 2 Output PFC TLV pkts, 0 Error PFC pkts
  0 PFC Pause Tx pkts, 0 Pause Rx pkts
  1 Input ETS Conf TLV Pkts, 1 Output ETS Conf TLV Pkts, 0 Error ETS Conf TLV
Pkts
  1 Input ETS Reco TLV pkts, 1 Output ETS Reco TLV pkts, 0 Error ETS Reco TLV
Pkts
```

Example of the `show interface DCBx detail` Command (Legacy CEE)

```
FTOS(conf-if-te-0/17-lldp)#do sho int te 1/14 dc d

E-ETS Configuration TLV enabled             e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled            r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled             p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled     f-Application Priority for FCOE
disabled
I-Application priority for iSCSI enabled    i-Application Priority for iSCSI
disabled
-------------------------------------------------------------------------

Interface TenGigabitEthernet 1/14
  Remote Mac Address 00:01:e8:8a:df:a0
    Port Role is Auto-Upstream
  DCBx Operational Status is Enabled
  Is Configuration Source? FALSE
  Local DCBx Compatibility mode is CEE
  Local DCBx Configured mode is CEE
  Peer Operating version is CEE
  Local DCBx TLVs Transmitted: ErPFi

Local DCBx Status
-----------------
  DCBx Operational Version is 0
  DCBx Max Version Supported is 0
  Sequence Number: 1
  Acknowledgment Number: 1
  Protocol State: In-Sync

Peer DCBx Status:
----------------
  DCBx Operational Version is 0
  DCBx Max Version Supported is 0
  Sequence Number: 1
  Acknowledgment Number: 1
  Total DCBx Frames transmitted 994
  Total DCBx Frames received 646
  Total DCBx Frame errors 0
  Total DCBx Frames unrecognized 0
```

The following table describes the `show interface DCBx detail` command fields.

Table 16. `show interface DCBx detail` Command Description

Field	Description
Interface	Interface type with chassis slot and port number.
Port-Role	Configured DCBx port role: auto-upstream, auto-downstream, config-source, or manual.
DCBx Operational Status	Operational status (enabled or disabled) used to elect a configuration source and internally propagate a DCB configuration. The DCBx operational status is the combination of PFC and ETS operational status.
Configuration Source	Specifies whether the port serves as the DCBx configuration source on the switch: true (yes) or false (no).
Local DCBx Compatibility mode	DCBx version accepted in a DCB configuration as compatible. In auto-upstream mode, a port can only received a DCBx version supported on the remote peer.
Local DCBx Configured mode	DCBx version configured on the port: CEE, CIN, IEEE v2.5, or Auto (port auto-configures to use the DCBx version received from a peer).
Peer Operating version	DCBx version that the peer uses to exchange DCB parameters.
Local DCBx TLVs Transmitted	Transmission status (enabled or disabled) of advertised DCB TLVs (see TLV code at the top of the show command output).
Local DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs.
Local DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs.
Local DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs.
Local DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs.
Local DCBx Status: Protocol State	Current operational state of DCBx protocol: ACK or IN-SYNC.
Peer DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs received from peer device.
Peer DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs received from peer device.
Peer DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs received from peer device.
Peer DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs received from peer device.
Total DCBx Frames transmitted	Number of DCBx frames sent from local port.

Field	Description
Total DCBx Frames received	Number of DCBx frames received from remote peer port.
Total DCBx Frame errors	Number of DCBx frames with errors received.
Total DCBx Frames unrecognized	Number of unrecognizable DCBx frames received.

PFC and ETS Configuration Examples

This section contains examples of how to configure and apply DCB input and output policies on an interface.

Using PFC and ETS to Manage Data Center Traffic

The following shows examples of using PFC and ETS to manage your data center traffic.

In the following example:

- Incoming SAN traffic is configured for priority-based flow control.
- Outbound LAN, IPC, and SAN traffic is mapped into three ETS priority groups and configured for enhanced traffic selection (bandwidth allocation and scheduling).
- One lossless queue is used.

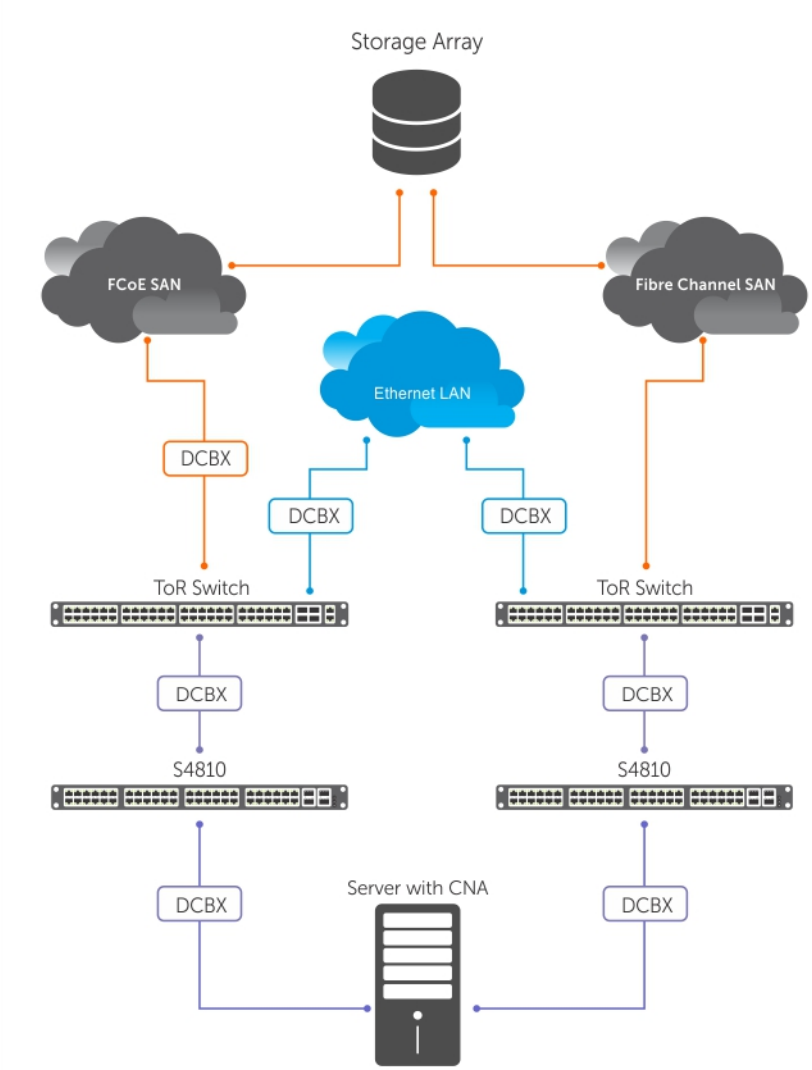

Figure 12. PFC and ETS Applied to LAN, IPC, and SAN Priority Traffic

QoS Traffic Classification: The `service-class dynamic dot1p` command has been used in Global Configuration mode to map ingress dot1p frames to the queues shown in the following table. For more information, refer to QoS dot1p Traffic Classification and Queue Assignment.

The following describes the dot1p-priority class group assignment

dot1p Value in the Incoming Frame	Priority Group Assignment
0	LAN
1	LAN
2	LAN
3	SAN

dot1p Value in the Incoming Frame	Priority Group Assignment
4	IPC
5	LAN
6	LAN
7	LAN

The following describes the priority group-bandwidth assignment.

Priority Group	Bandwidth Assignment
IPC	5%
SAN	50%
LAN	45%

PFC and ETS Configuration Command Examples

The following examples show PFC and ETS configuration commands to manage your data center traffic.

Example of Configuring QoS Priority-Queue Assignment to Honor Dot1p Priorities

```
FTOS(conf)# service-class dynamic dot1p

Or

FTOS(conf)# interface tengigabitethernet 0/1
FTOS(conf-if-te-0/1)# service-class dynamic dot1p
```

Example of Configuring a DCB Input Policy to Apply PFC to Lossless SAN Priority Traffic

```
FTOS(conf)# dcb-input ipc_san_lan
FTOS(conf-qos-policy-in)# pfc mode on
FTOS(conf-qos-policy-in)# pfc priority 3
```

Example of Configuring an ETS Priority Group

```
FTOS(conf)# priority-group san
FTOS(conf-pg)# priority-list 3
FTOS(conf-pg)# set-pgid 1
FTOS(conf-pg)# exit
FTOS(conf)# priority-group ipc
FTOS(conf-pg)# priority-list 4
FTOS(conf-pg)# set-pgid 2
FTOS(conf-pg)# exit
FTOS(conf)# priority-group lan
FTOS(conf-pg)# priority-list 0-2,5-7
FTOS(conf-pg)# set-pgid 3
FTOS(conf-pg)# exit
```

Example of Configuring an ETS Output Policy for Egress Traffic

```
FTOS(conf)# qos-policy-output san ets
FTOS(conf-qos-policy-out)# bandwidth-percentage 50
FTOS(conf-qos-policy-out)# exit
FTOS(conf)# qos-policy-output lan ets
FTOS(conf-qos-policy-out)# bandwidth-percentage 45
FTOS(conf-qos-policy-out)# exit
FTOS(conf)# qos-policy-output ipc ets
```

```
FTOS(conf-qos-policy-out)# bandwidth-percentage 5
FTOS(conf-qos-policy-out)# exit
```

Example of Configuring a DCB Output Policy to Apply ETS (Bandwidth Allocation and Scheduling) to IPC, SAN, and LAN Priority Traffic

```
FTOS(conf)# dcb-output ets
FTOS(conf-dcb-out)# priority-group san qos-policy san
FTOS(conf-dcb-out)# priority-group lan qos-policy lan
FTOS(conf-dcb-out)# priority-group ipc qos-policy ipc
```

Example of Applying DCB Input and Output Policies to an Interface

```
FTOS(conf)# interface tengigabitethernet 0/1
FTOS(conf-if-te-0/1)# dcb-policy input pfc
FTOS(conf-if-te-0/1)# dcb-policy output ets
```

Example of Configuring a QoS Output Policy to Specify Bandwidth Allocation to Different Traffic Types if DCBx Version is CIN

```
FTOS(conf)# qos-policy-output lan-q0
FTOS(conf-qos-policy-out)# bandwidth-percentage 20
FTOS(conf-qos-policy-out)# exit
FTOS(conf)#q os-policy-output lan-q3
FTOS(conf-qos-policy-out)# bandwidth-percentage 70
FTOS(conf-qos-policy-out)# exit
FTOS(conf)#policy-map-output ets-queues
```

Example of Creating a QoS Policy Map for DCBx CIN Bandwidth Allocation

```
FTOS(conf)# policy-map-output ets-queues
FTOS(conf-policy-map-out)# service-queue 0 qos-policy lan-q0
FTOS(conf-policy-map-out)# service-queue 3 qos-policy lan-q3
```

Example of Applying the QoS Policy Map for DCBx CIN Bandwidth Allocation to an Interface

```
FTOS(conf-if-te-0/1)# service-policy output ets-queues
```

Using PFC and ETS to Manage Converged Ethernet Traffic in a Switch Stack

The following example shows how to apply the DCB PFC input policy (`ipc_san_lan`) and ETS output policy (`ets`) on all ports in a switch stack.

This example references the PFC and ETS Configuration Examples section.

Example of Applying DCB PFC Input Policy and ETS Output Policy in a Switch Stack

```
FTOS(conf)# dcb-policy output stack-unit all stack-ports all ets
FTOS(conf)# dcb-policy input stack-unit all stack-ports all pfc
```

Hierarchical Scheduling in ETS Output Policies

ETS supports up to three levels of hierarchical scheduling.

For example, you can apply ETS output policies with the following configurations:

Priority group 1	Assigns traffic to one priority queue with 20% of the link bandwidth and strict-priority scheduling.
Priority group 2	Assigns traffic to one priority queue with 30% of the link bandwidth.
Priority group 3	Assigns traffic to two priority queues with 50% of the link bandwidth and strict-priority scheduling.

In this example, the configured ETS bandwidth allocation and scheduler behavior is as follows:

Unused bandwidth usage: Normally, if there is no traffic or unused bandwidth for a priority group, the bandwidth allocated to the group is distributed to the other priority groups according to the bandwidth percentage allocated to each group. However, when three priority groups with different bandwidth allocations are used on an interface:

- If priority group 3 has free bandwidth, it is distributed as follows: 20% of the free bandwidth to priority group 1 and 30% of the free bandwidth to priority group 2.
- If priority group 1 or 2 has free bandwidth, (20 + 30)% of the free bandwidth is distributed to priority group 3. Priority groups 1 and 2 retain whatever free bandwidth remains up to the (20+ 30)%.

Strict-priority groups: If two priority groups have strict-priority scheduling, traffic assigned from the priority group with the higher priority-queue number is scheduled first. However, when three priority groups are used and two groups have strict-priority scheduling (such as groups 1 and 3 in the example), the strict priority group whose traffic is mapped to one queue takes precedence over the strict priority group whose traffic is mapped to two queues.

Therefore, in this example, scheduling traffic to priority group 1 (mapped to one strict-priority queue) takes precedence over scheduling traffic to priority group 3 (mapped to two strict-priority queues).

39

Fibre Channel over Ethernet for FC Flex IO Modules

FCoE provides a converged Ethernet network that allows the combination of storage-area network (SAN) and LAN traffic on a Layer 2 link by encapsulating Fibre Channel data into Ethernet frames.

The Fibre Channel (FC) Flex IO module is supported on Dell Networking MXL 10/40GbE Switch and Dell PowerEdge IO Aggregator (IOA). The MXL and IOA switch installed with the FC Flex IO module functions as a top-of-rack edge switch that supports converged enhanced ethernet (CEE) traffic — Fibre channel over Ethernet (FCoE) for storage, inter-process communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics.

FCoE works with the Ethernet enhancements provided in data center bridging (DCB) to support lossless (no-drop) SAN and LAN traffic. In addition, DCB provides flexible bandwidth sharing for different traffic types, such as LAN and SAN, according to 802.1p priority classes of service. DCBx should be enabled on the system before the FIP snooping feature is enabled.

All of the commands that are supported for FCoE on the MXL and I/O Aggregator apply to the FC Flex IO modules. Similarly, all of the configuration procedures and the settings that are applicable for FCoE on the MXL and I/O Aggregator are valid for the FC Flex IO modules.

Legacy FIP snooping-related commands are supported on 10G optional modules when the Fiber channel capability is disabled (by using the `no feature fc` command). If the FC capability is enabled, only the `fip-snooping max-sessions-per-enodemac` command is supported and the remaining FIP snooping commands are not supported.

For details on the command syntaxes and the keywords, refer the *FCoE chapter of the MXL and I/O Aggregator Command Reference Guides of Release 9.2.0.0*. For details on the step-wise instructions to enable and configure FCoE settings, refer the *FCoE chapter of the MXL and I/O Aggregator Configuration Guides of Release 9.2.0.0*.

40

NPIV Proxy Gateway for FC Flex IO Modules

The N-port identifier virtualization (NPIV) Proxy Gateway (NPG) feature provides FCoE-FC bridging capability on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module switch, allowing server CNAs to communicate with SAN fabrics over the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

To configure the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module to operate as an NPIV proxy gateway, use the following commands:

dcb-map

Create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic. Apply the DCB map to an Ethernet interface.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`dcb-map` *`map-name`*	
Parameters	***map-name***	Enter a DCB map name. The maximum number of alphanumeric characters is 32.
Defaults	None. There are no pre-configured PFC and ETS settings on M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module Ethernet interfaces. With auto-detection of DCB enabled, a DCB map named 'dcb-map' is applied on all the Ethernet interfaces on which the DCBx frames are observed.	
Command Modes	CONFIGURATION INTERFACE	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.
Usage Information	A DCB map is a template used to configure DCB parameters and apply them on converged Ethernet interfaces. DCB parameters include priority-based flow control (PFC) and enhanced traffic selection (ETS). To display the PFC and ETS settings in DCB maps, enter the `show qos dcb-map` command. Use the `dcb-map` command to create a DCB map to specify PFC and ETS settings and apply it on Ethernet ports. After you apply a DCB map to an interface, the PFC and ETS settings in the map are applied when the Ethernet port is enabled. DCBx is enabled on Ethernet ports by default. The `dcb-map` command is supported only on physical Ethernet interfaces.	

	To remove a DCB map from an interface, enter the `no dcb-map map-name` command in Interface configuration mode.
Related Commands	show qos dcb-map– displays the dcb-map profiles configured on the system. dcb-map stack-unit all stack-ports all– applies a DCB map on all ports of a switch stack.

description (for FCoE maps)

In an FCoE map, add a text description of the FCoE and FC parameters used to transmit storage traffic over an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway in a converged fabric.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`description text`	
Parameters	***text***	Enter a maximum of 32 characters.
Defaults	None	
Command Modes	FCOE MAP	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.
Usage Information	The text description is displayed in `show fcoe-map` command output.	
Related Commands	fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric. show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.	

fabric

Apply an FCoE map on a fabric-facing Fibre Channel (FC) port.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`fabric map-name`	
Parameters	***map-name***	Maximum: 32 alphanumeric characters.
Defaults	None	

Command Modes	INTERFACE FIBRE_CHANNEL	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

An FCoE map is a template used to map FCoE and FC parameters in a converged fabric. An FCoE map is used to virtualize upstream FC ports on an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway so that they appear to downstream server CNA ports as FCoE forwarder (FCF) ports on an FCoE network. When applied to FC and Ethernet ports on an NPIV proxy gateway, an FCoE map allows the switch to operate as an FCoE-FC bridge between an FC SAN and an FCoE network by providing FCoE-enabled servers and switches with the necessary parameters to log in to a SAN fabric. Use the `fcoe-map` command to create an FCoE map.

On an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway, you cannot apply an FCoE map on fabric-facing FC ports and server-facing Ethernet ports.

After you apply an FCoE map on an FC interface, when the port is enabled (`no shutdown`), the NPIV proxy gateway starts sending FIP multicast advertisements on behalf of the FC port to downstream servers in order to advertise the availability of a new FCF port on the FCoE VLAN.

To remove an FCoE map from an FC interface, enter the `no fabric` *`map-name`* command in Interface configuration mode.

Related Commands

fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric.

interface fibrechannel — configures a Fibre Channel interface on the switch. .

show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.

fabric-id vlan

In an FCoE map, configure the association between the dedicated VLAN used to carry FCoE traffic between servers and a SAN, and the fabric where the desired storage arrays are installed.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax `fabric-id` *`fabric-num`* `vlan` *`vlan-id`*

Parameters

fabric-id ***fabric-num***	Enter a fabric ID number that is the same as the ID number of the dedicated VLAN used to carry FCoE storage traffic to the fabric specified in the FCoE map. You can enter a fabric ID in the range 1–4094.
vlan*vlan-id*	Enter the ID number of the dedicated VLAN used to carry FCoE storage traffic between servers and a SAN fabric and specified with the `vlan` command in the FCoE map.

Defaults	None
Command Modes	FCOE MAP
Command History	**Version 9.3.0.0** Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.
Usage Information	In the `fabric-id vlan` command, the fabric and VLAN ID numbers must be the same. In each FCoE map, the fabric ID, FC-MAP value, and FCoE VLAN parameters must be unique. To remove a fabric-VLAN association from an FCoE map, enter the `no fabric-id vlan` command. You must first create a VLAN and then specify the configured VLAN ID in the `fabric-id vlan` command. Otherwise, the following error message is displayed. `FTOS(conf-fcoe-f)#fabric-id 10 vlan 10 % Error: Vlan 10 does not exist`
Related Commands	fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric. show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.

fcf-priority

In an FCoE map, configure the priority used by a server CNA to select an upstream FCoE forwarder (FCF).

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`fcf—priority` *priority*
Parameters	***priority*** Enter the priority assigned to the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway, which appears to a downstream server CNA as an FCF. The range of FCF priority values is from 1 to 255.
Defaults	128
Command Modes	FCOE MAP
Command History	**Version 9.3.0.0** Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.
Usage Information	The FCF priority you assign to an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module is used by server CNAs to select an upstream FCF to use for a fabric login (FLOGI).

To remove a configured FCF priority from an FCoE map, enter the `no fcf-priority` command.

Related Commands

fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric.

show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.

fc-map

In an FCoE map, configure the FCoE mapped address prefix (FC-MAP) value which is used to identify FCoE traffic transmitted on the FCoE VLAN for the specified fabric.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax `fc—map fc-map-value`

Parameters

fc-map-value	Enter the unique MAC address prefix used by a SAN fabric. The range of FC-MAP values is from 0EFC00 to 0EFCFF.

Defaults None

Command Modes FCOE MAP

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

The FC-MAP value you enter must match the FC-MAP value used by an FC switch or FCoE forwarder (FCF) in the fabric. An FCF switch accepts only FCoE traffic that uses the correct FC-MAP value.

The FC-MAP value is used to generate the fabric-provided MAC address (FP-MAC). The FPMA is used by servers to transmit FCoE traffic to the fabric. An FC-MAP can be associated with only one FCoE VLAN and vice versa.

In an FCoE map, the FC-MAP value, fabric ID, and FCoE VLAN parameters must be unique.

To remove a configured FC-MAP value from an FCoE map, enter the `no fc-map` command.

Related Commands

fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric.

show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.

fcoe priority-bits

Configure the FCoE priority advertised for the FCoE protocol in application priority TLVs.

Syntax	`fcoe priority-bits` *`priority-bitmap`* To remove the configured FCoE priority, use the `no fcoe priority-bits` command.	
Parameters	***priority-bitmap***	Enter the priority-bitmap range. The range is from 1 to FF.
Defaults	0x8	
Command Modes	PROTOCOL LLDP	
Command History	**Version 9.3.0.0**	Introduced on the FC Flex IO module installed in the M I/O Aggregator and MXL 10/40GbE Switch
Usage Information	This command is available at the global level only.	

fcoe-map

Create an FCoE map which contains the parameters used to configure the links between server CNAs and a SAN fabric. Apply the FCoE map on a server-facing Ethernet port.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`fcoe-map` *`map-name`*	
Parameters	***map-name***	Maximum: 32 alphanumeric characters.
Defaults	None on the MXL 10/40GbE Switch with FC Flex IO modules On the I/O Aggregator with FC Flex IO modules, the following parameters are applied on all the FC Flex IO module interfaces: • Description: SAN_FABRIC • Fabric-id: 1002 • Fcoe-vlan: 1002 • Fc-map: 0x0efc00 • Fcf-priority: 128 • Fka-adv-period: 8000mSec • Keepalive: enable • Vlan priority: 3	

Command Modes	CONFIGURATION INTERFACE
Command History	**Version 9.3.0.0** Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.
Usage Information	An FCoE map is a template used to map FCoE and FC parameters in a converged fabric. An FCoE map is used to virtualize upstream FC ports on an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway so that they appear to downstream server CNA ports as FCoE forwarder (FCF) ports on an FCoE network. When applied to FC and Ethernet ports on an NPIV proxy gateway, an FCoE map allows the switch to operate as an FCoE-FC bridge between an FC SAN and an FCoE network by providing FCoE-enabled servers and switches with the necessary parameters to log in to a SAN fabric. On an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway, you cannot apply an FCoE map is applied on fabric-facing FC ports and server-facing 10–Gigabit Ethernet ports. An FCoE map consists of the following parameters: the dedicated FCoE VLAN used for storage traffic, the destination SAN fabric (FC-MAP value), FCF priority used by a server, and the FIP keepalive (FKA) advertisement timeout. In each FCoE map, the fabric ID, FC-MAP value, and FCoE VLAN parameters must be unique. Use one FCoE map to access one SAN fabric. You cannot use the same FCoE map to access different fabrics. To remove an FCoE map from an Ethernet interface, enter the `no fcoe-map` *`map-name`* command in Interface configuration mode.
Related Commands	interface fibrechannel — configures a Fibre Channel interface on the switch. show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.

fka-adv-period

In an FCoE map, configure the time interval used to transmit FIP keepalive (FKA) advertisements.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`fka-adv-period` *`seconds`*
Parameters	***seconds*** Enter the time period (in seconds) used to send FIP keepalive messages to peer devices. The range is from 8 to 90 seconds.
Defaults	8 seconds
Command Modes	FCOE MAP

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

To delete the FIP keepalive time period from an FCoE map, enter the `no fka-adv-erpiod` command.

Related Commands

fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric.

show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.

interface vlan (NPIV proxy gateway)

Create a dedicated VLAN to be used to send and receive Fibre Channel traffic over FCoE links between servers and a fabric over an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax

```
interface vlan vlan-id
```

Parameters

vlan-id	Enter a number as the VLAN Identifier. The range is 1 to 4094.

Defaults

Not configured.

Command Modes

CONFIGURATION

Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module configured as an NPIV proxy gateway.

Usage Information

FCoE storage traffic received from servers on an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway is de-capsulated into Fibre Channel packets and forwarded over FC links to SAN switches in a specified fabric. You must configure a separate FCoE VLAN for each fabric to which FCoE traffic is forwarded. Any non-FCoE traffic sent on a dedicated FCoE VLAN will be dropped.

You configure the association between a dedicated VLAN, which carries FCoE traffic from server CNAs over the NPIV proxy gateway to a SAN fabric in which destination storage arrays are installed, in an FCoE map by using the `fabric id vlan` command.

When you apply an FCoE map to a server-facing Ethernet port, the port is automatically configured as a tagged member of the FCoE VLAN.

For more information about VLANs and the commands to configure them, refer to the Virtual LAN (VLAN) Commands section of the Layer 2 chapter.

Example (Single Range)	`FTOS(conf)#interface vlan 10` `FTOS(conf-if-vl-3)#`
Related Commands	fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric. show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.

keepalive

In an FCoE map, enable the monitoring of FIP keepalive messages (if it is disabled).

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`keepalive`	
Parameters	None	
Defaults	FIP keepalive monitoring is enabled on Ethernet and Fibre Channel interfaces.	
Command Modes	FCOE MAP	
Command History	**Version 9.3.0.0**	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.
Usage Information	FIP keepalive (FKA) messaging is used to detect if other FCoE devices are reachable. To remove FIP keepalive monitoring from an FCoE map, enter the `no keepalive` command.	
Related Commands	fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric. show fcoe-map— displays the Fibre Channel and FCoE configuration parameters in FCoE maps.	

priority-group bandwidth pfc

Configure the ETS bandwidth allocation and PFC mode used to manage port traffic in an 802.1p priority group.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax	`priority-group` *`group-num`* `{bandwidth` *`percentage`*`| strict-priority} pfc {on | off}`	
Parameters	**priority-group** ***group-num***	Enter the keyword `priority-group` followed by the number of an 802.1p priority group. Use the `priority-pgid` command to create the priority groups in a DCB map.

bandwidth *percentage*	Enter the keyword `bandwidth` followed by a bandwidth percentage allocated to the priority group. The range of valid values is 1 to 100. The sum of all allocated bandwidth percentages in priority groups in a DCB map must be 100%.
strict-priority	Configure the priority-group traffic to be handled with strict priority scheduling. Strict-priority traffic is serviced first, before bandwidth allocated to other priority groups is made available.
pfc {on | off}	Configure whether priority-based flow control is enabled (on) or disabled (off) for port traffic in the priority group.

Defaults

None

Command Modes

DCB MAP

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

Use the `dcb-map` command to configure priority groups with PFC and/or ETS settings and apply them to Ethernet interfaces.

Use the `priority-pgid` command to map 802.1p priorities to a priority group. You can assign each 802.1p priority to only one priority group. A priority group consists of 802.1p priority values that are grouped together for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.

Repeat the `priority-group bandwidth pfc` command to configure PFC and ETS traffic handling for each priority group in a DCB map.

You can enable PFC on a maximum of two priority queues.

If you configure more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.

To remove a priority-group configuration in a DCB map, enter the `no priority-group bandwidth pfc` command.

By default, equal bandwidth is assigned to each dot1p priority in a priority group. Use the `bandwidth` parameter to configure the bandwidth percentage assigned to a priority group. The sum of the bandwidth allocated to all priority groups in a DCB map must be 100% of the bandwidth on the link. You must allocate at least 1% of the total port bandwidth to each priority group.

Related Commands

dcb-map – creates a DCB map to configure PFC and ETS parameters and applies the PFC and ETS settings on Ethernet ports.

priority-pgid – configures the 802.1p priority traffic in a priority group for a DCB map.

show fcoe-map

Display the Fibre Channel and FCoE configuration parameters in FCoE maps.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax

```
show fcoe-map [brief | map-name]
```

Parameters

brief	Displays an overview of currently configured FCoE maps.
map-name	Displays the FC and FCoE configuration parameters in a specified FCoE map. The FCoE map is applied on Ethernet (FCoE) and FC ports to transmit FC storage traffic to a specified fabric.

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

Use the `show fcoe-map` command to display the FC and FCoE parameters used to configure server-facing Ethernet (FCoE) and fabric-facing FC ports in all FCoE maps on an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway.

In each FCoE map, the values for the fabric ID and FC-MAP that identify the SAN fabric to which FC storage traffic is sent, and the FCoE VLAN to be used must be unique.

An FCoE map is used to identify the SAN fabric to which FCoE storage traffic is sent and to virtualize M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module FC ports so that they appear to downstream server CNA ports as FCoE Forwarder (FCF) ports on an FCoE network.

The following table describes the `show fcoe-map brief` output shown in the example below.

Field	Description
Fabric-Name	Name of a SAN fabric.
Fabric ID	The ID number of the SAN fabric to which FC traffic is forwarded.
VLAN ID	The dedicated FCoE VLAN used to transport FCoE storage traffic between servers and a fabric over the NPIV proxy gateway. The configured VLAN ID must be the same as the fabric ID.

Field	Description
FC-MAP	FCoE MAC address-prefix value - The unique 24-bit MAC address prefix that identifies a fabric.
FCF Priority	The priority used by a server to select an upstream FCoE forwarder.
Config-State	Indicates whether the configured FCoE and FC parameters in the FCoE map are valid: Active (all mandatory FCoE and FC parameters are correctly configured) or Incomplete (either the FC-MAP value, fabric ID, or VLAN ID are not correctly configured).
Oper-State	Operational status of link to the fabric: Up (link is up and transmitting FC traffic), Down (link is down and not transmitting FC traffic), Link-wait (link is up and waiting for FLOGI to complete on peer FC port), or Removed (port has been shut down).

The following table describes the `show fcoe-map` *`map-name`* output shown in the example below.

Field	Description
Fabric-Name	Name of a SAN fabric.
Fabric ID	The ID number of the SAN fabric to which FC traffic is forwarded.
VLAN ID	The dedicated FCoE VLAN used to transport FCoE storage traffic between servers and a fabric over the NPIV proxy gateway. The configured VLAN ID must be the same as the fabric ID.
VLAN priority	FCoE traffic uses VLAN priority 3. (This setting is not user-configurable.)
FC-MAP	FCoE MAC address-prefix value - The unique 24-bit MAC address prefix that identifies a fabric.
FKA-ADV-period	Time interval (in seconds) used to transmit FIP keepalive advertisements.
FCF Priority	The priority used by a server to select an upstream FCoE forwarder.
Config-State	Indicates whether the configured FCoE and FC parameters in the FCoE map are valid: Active (all mandatory FCoE and FC parameters are correctly configured) or Incomplete (either the FC-MAP value, fabric ID, or VLAN ID are not correctly configured).
Oper-State	Operational status of link to the fabric: Up (link is up and transmitting FC traffic), Down (link is down and not transmitting FC traffic), Link-wait (link is up and waiting for

Field	Description
	FLOGI to complete on peer FC port), or Removed (port has been shut down).
Members	M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module Ethernet and FC ports that are members of the dedicated FCoE VLAN that carries storage traffic to the specified fabric.

Example

```
FTOS#show fcoe-map brief
Fabric-Name  Fabric-Id  Vlan-Id  FC-MAP  FCF-Priority Config-
State Oper-State
 test        16          16      0efc02  128
ACTIVE        UP
 cnatest     1003           1003    0efc03  128
ACTIVE        UP
 sitest      1004        1004    0efc04  128
ACTIVE        DOWN


FTOS#show fcoe-map si

Fabric Name         si
Fabric Id           1004
Vlan Id             1004
Vlan priority       3
FC-MAP              0efc04
FKA-ADV-Period      8
Fcf Priority        128
Config-State        ACTIVE
Oper-State          DOWN
Members
```

Related Commands

fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric.

show npiv devices

Display the FCoE and FC devices currently logged in to an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway.

M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module

Syntax

```
show npiv devices [brief]
```

Parameters

brief	Displays an overview of current server CNA-fabric connections over an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway.

Command Modes

- EXEC
- EXEC Privilege

Command History

Version 9.3.0.0	Introduced on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

Usage Information

Use the `show npiv devices` command to display information on the server CNA, server-facing Ethernet and fabric-facing FC ports, and the SAN fabric in each server-fabric connection over an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module NPIV proxy gateway.

The following table describes the `show npiv devices brief` output shown in the example below.

Field	Description
ENode-Intf	M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module Ethernet interface (*slot/port*) to which a server CNA is connected.
ENode-WWPN	Worldwide port name (WWPN) of a server CNA port.
FCoE-Vlan	VLAN ID of the dedicated VLAN used to transmit FCoE traffic to and from the fabric.
Fabric-Intf	Fabric-facing Fibre Channel port (*slot/port*) on which FC traffic is transmitted to the specified fabric.
Fabric-Map	Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.
LoginMethod	Method used by the server CNA to log in to the fabric; for example: FLOGI - ENode logged in using a fabric login (FLOGI). FDISC - ENode logged in using a fabric discovery (FDISC).
Status	Operational status of the link between a server CNA port and a SAN fabric: Logged In - Server has logged in to the fabric and is able to transmit FCoE traffic.

Example

```
FTOS# show npiv devices brief

Total NPIV Devices = 2
-------------------------------------------------------------------
---------------------------

ENode-Intf  ENode-WWPN           FCoE-Vlan  Fabric-Intf
Fabric-Map   LoginMethod  Status
-------------------------------------------------------------------
---------------------------

Te 0/12     20:01:00:10:18:f1:94:20   1003      Fc 0/5
fid_1003     FLOGI       LOGGED_IN
Te 0/13     10:00:00:00:c9:d9:9c:cb   1003      Fc 0/0
fid_1003     FDISC       LOGGED_IN
```

Usage Information

The following table describes the `show npiv devices` output shown in the example below.

Field	Description
ENode [*number*]	A server CNA that has successfully logged in to a fabric over an M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module Ethernet port in ENode mode.
Enode MAC	MAC address of a server CNA port.
Enode Intf	Port number of a server-facing Ethernet port operating in ENode mode.
FCF MAC	Fibre Channel forwarder MAC: MAC address of M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module FCF interface.
Fabric Intf	Fabric-facing Fibre Channel port (*slot/port*) on which FCoE traffic is transmitted to the specified fabric.
FCoE VLAN	ID of the dedicated VLAN used to transmit FCoE traffic from a server CNA to a fabric and configured on both the server-facing M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module port and server CNA port.
Fabric Map	Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.
Enode WWPN	Worldwide port name of the server CNA port.
Enode WWNN	Worldwide node name of the server CNA.
FCoE MAC	Fabric-provided MAC address (FPMA). The FPMA consists of the FC-MAP value in the FCoE map and the FC-ID provided by the fabric after a successful FLOGI. In the FPMA, the most significant bytes are the FC-MAP; the least significant bytes are the FC-ID.
FC-ID	FC port ID provided by the fabric.
LoginMethod	Method used by the server CNA to log in to the fabric; for example, FLOGI or FDISC.
Secs	Number of seconds that the fabric connection is up.
State	Status of the fabric connection: logged in.

Example

```
ENode[0]:
ENode MAC    :    00:10:18:f1:94:21
ENode Intf   :    Te 0/12
FCF MAC      :    5c:f9:dd:ef:10:c8
Fabric Intf  :    Fc 0/5
FCoE Vlan    :    1003
Fabric Map   :    fid_1003
ENode WWPN   :    20:01:00:10:18:f1:94:20
ENode WWNN   :    20:00:00:10:18:f1:94:21
FCoE MAC     :    0e:fc:03:01:02:01
FC-ID        :    01:02:01
```

```
LoginMethod  :     FLOGI
Secs         :     5593
Status       :     LOGGED_IN

ENode[1]:
ENode MAC    :     00:10:18:f1:94:22
ENode Intf   :     Te 0/13
FCF MAC      :     5c:f9:dd:ef:10:c9
Fabric Intf  :     Fc 0/0
FCoE Vlan    :     1003
Fabric Map   :     fid_1003
ENode WWPN   :     10:00:00:00:c9:d9:9c:cb
ENode WWNN   :     10:00:00:00:c9:d9:9c:cd
FCoE MAC     :     0e:fc:03:01:02:02
FC-ID        :     01:02:01
LoginMethod  :     FDISC
Secs         :     5593
Status       :     LOGGED_IN
```

Related Commands

dcb-map — creates a DCB map to configure DCB parameters on Ethernet ports that support converged Ethernet traffic.

fcoe-map — creates an FCoE map which contains the parameters used in the communication between servers and a SAN fabric.

NPIV Proxy Gateway Configuration on FC Flex IO Modules

The Fibre Channel (FC) Flex IO module is supported on Dell Networking MXL 10/40GbE Switch and Dell PowerEdge IO Aggregator (IOA). The MXL and IOA switch installed with the FC Flex IO module functions as a top-of-rack edge switch that supports converged enhanced ethernet (CEE) traffic — Fibre channel over Ethernet (FCoE) for storage, inter-process communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics.

The N-port identifier virtualization (NPIV) proxy gateway (NPG) provides FCoE-FC bridging capability on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module.

This chapter describes how to configure and use an NPIV proxy gateway on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module in a storage area network (SAN).

NPIV Proxy Gateway Operations and Capabilities

Benefits of an NPIV Proxy Gateway

The MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module functions as a top-of-rack edge switch that supports Converged Enhanced Ethernet (CEE) traffic — FCoE for storage, Inter-process Communication (IPC) for servers, and Ethernet LAN (IP cloud) for data — as well as Fibre Channel (FC) links to one or more SAN fabrics.

Using an NPIV proxy gateway (NPG) helps resolve the following problems in a storage area network:

- Fibre Channel storage networks typically consist of servers connected to edge switches, which are connected to SAN core switches. As the SAN grows, it is necessary to add more ports and SAN switches. This results in an increase in the required domain IDs, which may surpass the upper limit of 239 domain IDs supported in the SAN network. An NPG avoids the need for additional domain IDs because it is deployed outside the SAN and uses the domain IDs of core switches in its FCoE links.
- With the introduction of 10GbE links, FCoE is being implemented for server connections to optimize performance. However, a SAN traditionally uses Fibre Channel to transmit storage traffic. FCoE servers

require an efficient and scalable bridging feature in order to access FC storage arrays, which an NPG provides.

NPIV Proxy Gateway Operation

Consider a sample scenario of NPG operation. An M1000e chassis configured as an NPG does not join a SAN fabric, but functions as an FCoE-FC bridge that forwards storage traffic between servers and core SAN switches. The core switches forward SAN traffic to and from FC storage arrays.

An M1000e chassis FC port is configured as an N (node) port that logs in to an F (fabric) port on the upstream FC core switch and creates a channel for N-port identifier virtualization. NPIV allows multiple N-port fabric logins at the same time on a single, physical Fibre Channel link.

Converged Network Adapter (CNA) ports on servers connect to the M1000e chassis Ten-Gigabit Ethernet ports and log in to an upstream FC core switch through the FC Flex IO module N port. Server fabric login (FLOGI) requests are converted into fabric discovery (FDISC) requests before being forwarded by the FC Flex IO module to the FC core switch.

Servers use CNA ports to connect over FCoE to an Ethernet port in ENode mode on the NPIV proxy gateway. FCoE transit with FIP snooping is automatically enabled and configured on the M1000e gateway to prevent unauthorized access and data transmission to the SAN network (see FCoE Transit). FIP is used by server CNAs to discover an FCoE switch operating as an FCoE forwarder (FCF).

The NPIV proxy gateway aggregates multiple locally connected server CNA ports into one or more upstream N port links, conserving the number of ports required on an upstream FC core switch while providing an FCoE-to-FC bridging functionality. The upstream N ports on an M1000e can connect to the same or multiple fabrics.

Using an FCoE map applied to downstream (server-facing) Ethernet ports and upstream (fabric-facing) FC ports, you can configure the association between a SAN fabric and the FCoE VLAN that connects servers over the NPIV proxy gateway to FC switches in the fabric. An FCoE map virtualizes the upstream SAN fabric as an FCF to downstream CNA ports on FCoE-enabled servers as follows:

- As soon as an FC N port comes online (**no shutdown** command), the NPG starts sending FIP multicast advertisements, which contain the fabric name derived from the 64-bit worldwide name (WWN) of the principal SAN switch. (The principal switch in a fabric is the FC switch with the lowest domain ID.)
- When you apply the FCoE map to a server-facing Ethernet port in ENode mode, ACLs are automatically configured to allow only FCoE traffic from servers that perform a successful FLOGI on the FC switch. All other traffic on the VLAN is denied.

You can specify one or more upstream N ports in an FCoE map. The FCoE map also contains the VLAN ID of the dedicated VLAN used to transmit FCoE traffic between the SAN fabric and servers.

NPIV Proxy Gateway: Protocol Services

An MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG provides the following protocol services:

- Fibre Channel service to create N ports and log in to an upstream FC switch.
- FCOE service to perform:
- Virtualization of FC N ports on an NPG so that they appear as FCOE FCFs to downstream servers.
- NPIV service to perform the association and aggregation of FCoE servers to upstream F ports on core switches (through N ports on the NPG). Conversion of server FLOGIs and FDISCs, which are received

over MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module ENode ports, are converted into FDISCs addressed to the upstream F ports on core switches.

NPIV Proxy Gateway Functionality

An MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG provides the following functionality in a storage area network:

- FIP Snooping bridge that provides security for FCoE traffic using ACLs (see FCoE Transit chapter).
- FCoE gateway that provides FCoE-to-FC bridging. N-port virtualization using FCoE maps exposes upstream F ports as FCF ports to downstream server-facing ENode ports on the NPG (see FCoE Maps).

NPIV Proxy Gateway: Terms and Definitions

The following table describes the terms used in an NPG configuration on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module.

Table 17. MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPIV Proxy Gateway: Terms and Definitions

Term	Description
FC port	Fibre Channel port on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module FC module that operates in autosensing, 2, 4, or 8 Gigabit mode. On an NPIV proxy gateway, an FC port can be used as a downlink for a server connection and an uplink for a fabric connection.
F port	Port mode of an FC port connected to an end node (N) port on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPIV proxy gateway.
N port	Port mode of an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module FC port that connects to an F port on an FC switch in a SAN fabric. On an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPIV proxy gateway, an N port also functions as a proxy for multiple server CNA-port connections
ENode port	Port mode of a server-facing MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet port that provides access to FCF functionality on a fabric.
CNA port	N-port functionality on an FCoE-enabled server port. A converged network adapter (CNA) can use one or more Ethernet ports. CNAs have the capability to encapsulate Fibre Channel frames in Ethernet for FCoE transport and de-encapsulate Fibre Channel frames from FCoE to native Fibre Channel.
DCB map	Template used to configure DCB parameters, including priority-based flow control (PFC) and enhanced transmission selection (ETS), on CEE ports.
Fibre Channel fabric	Network of Fibre Channel devices and storage arrays that inter-operate and communicate.
FCF	Fibre Channel forwarder: FCoE-enabled switch that can forward FC traffic to both downstream FCoE and upstream FC devices. An NPIV proxy gateway functions as an FCF to export upstream F port configurations to downstream server CNA ports.

Term	Description
FC-MAP	FCoE MAC-address prefix - The unique 24-bit MAC address prefix in FCoE packets used to generate a fabric-provided MAC address (FPMA). The FPMA is required to send FCoE packets from a server to a SAN fabric.
FCoE map	Template used to configure FCoE and FC parameters on Ethernet and FC ports in a converged fabric.
FCoE VLAN	VLAN dedicated to carrying only FCoE traffic between server CNA ports and a SAN fabric. (FCoE traffic must travel in a VLAN.) When you apply an FCoE map on a port, FCoE is enabled on the port. All non-FCoE traffic is dropped on an FCoE VLAN.
FIP	FCoE Initialization Protocol: Layer 2 protocol for endpoint discovery, fabric login, and fabric association. FIP is used by server CNAs to discover an upstream FCoE switch operating as an FCF. FIP keepalive messages maintain the connection between an FCoE initiator and an FCF.
NPIV	N-port identifier virtualization: The capability to map multiple FCoE links from downstream ports to a single upstream FC link.
principal switch	The switch in a fabric with the lowest domain number. The principal switch accesses the master name database and the zone/zone set database.

DCB Maps

A Data Center Bridging (DCB) map is used to configure DCB functionality, such as PFC and ETS, on MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet ports that support CEE traffic and are DCBx-enabled by default. For more information, on PFC and ETS, see Data Center Bridging (DCB).

By default, no PFC and ETS settings in a DCB map are applied to MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet ports when they are enabled. On an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG, you must configure PFC and ETS parameters in a DCB map and then apply the map to server-facing Ethernet ports (see the “Creating a DCB map" section).

FCoE Maps

An FCoE map is used to identify the SAN fabric to which FCoE storage traffic is sent. Using an FCoE map, an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG operates as an FCoE-FC bridge between an FC SAN and FCoE network by providing FCoE-enabled servers and switches with the necessary parameters to log in to a SAN fabric.

An FCoE map applies the following parameters on server-facing Ethernet and fabric-facing FC ports on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module:

- The dedicated FCoE VLAN used to transport FCoE storage traffic.
- The FC-MAP value used to generate a fabric-provided MAC address.
- The association between the FCoE VLAN ID and FC fabric ID where the desired storage arrays are installed. Each Fibre Channel fabric serves as an isolated SAN topology within the same physical network.
- The priority used by a server to select an upstream FCoE forwarder (FCF priority)
- FIP keepalive (FKA) advertisement timeout

NOTE:
In each FCoE map, the fabric ID, FC-MAP value, and FCoE VLAN must be unique. Use one FCoE map to access one SAN fabric. You cannot use the same FCoE map to access different fabrics.

When you configure an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module as an NPG, FCoE transit with FIP snooping is automatically enabled and configured using the parameters in the FCoE map applied to server-facing Ethernet and fabric-facing FC interfaces (see FIP Snooping on an NPIV Proxy Gateway).

After you apply an FCoE map on an FC port, when you enable the port (no shutdown), the NPG starts sending FIP multicast advertisements on behalf of the FC port to downstream servers in order to advertise the availability of a new FCF port on the FCoE VLAN. The FIP advertisement also contains a keepalive message to maintain connectivity between a SAN fabric and downstream servers.

Configuring an NPIV Proxy Gateway

Prerequisite: Before you configure an NPIV proxy gateway on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module:

DCB is enabled by default on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module. Auto-negotiated DCBx is enabled for converged traffic by default on all MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet ports. FCoE transit with FIP snooping is automatically enabled when you configure Fibre Channel capability on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module (see Enabling Fibre Channel Capability on the Switch).

To configure NPG operation on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module, follow these general configuration steps:

1. Enabling Fibre Channel Capability on the Switch
2. Creating a DCB map.
3. Applying a DCB map on server-facing Ethernet ports.
4. Creating an FCoE VLAN.
5. Creating an FCoE map.
6. Applying an FCoE map on server-facing Ethernet ports.
7. Applying an FCoE Map on fabric-facing FC ports.

Enabling Fibre Channel Capability on the Switch

You must first enable an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module that you want to configure as an NPG for the Fibre Channel protocol. When you enable Fibre Channel capability, FCoE transit with FIP snooping is automatically enabled on all VLANs on the switch, using the default FCoE transit settings. For more information, see Enabling FIP Snooping on VLANs.

Task	Command	Command Mode
Enable an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module for the Fibre Channel protocol.	`feature fc`	CONFIGURATION

Creating a DCB map

Configure the priority-based flow control (PFC) and enhanced traffic selection (ETS) settings in a DCB map before you can apply them on downstream server-facing ports on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG.

Step	Task	Command	Command Mode
1	Create a DCB map to specify PFC and ETS settings for groups of dot1p priorities.	dcb-map *name*	CONFIGURATION
2	Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group or whether priority group traffic should be handled with strict priority scheduling. The sum of all allocated bandwidth percentages must be 100%. Strict-priority traffic is serviced first. Afterwards, bandwidth allocated to other priority groups is made available and allocated according to the specified percentages. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups. **Restriction**: You can enable PFC on a maximum of two priority queues. Repeat this step to configure PFC and ETS traffic handling for each priority group; for example: **priority-group 0 bandwidth 60 pfc off priority-group 1 bandwidth 20 pfc on priority-group 2 bandwidth 20 pfc on priority-group 4 strict-priority pfc off**	priority-group *group_num* {**bandwidth** *percentage* | **strict-priority**} **pfc** {**on** | **off**}	DCB MAP
3	Specify the priority group ID number to handle VLAN traffic for each dot1p class-of-service: 0 through 7. Leave a space between each priority group number. For example: **priority-pgid 0 0 0 1 2 4 4 4** Where dot1p priorities 0, 1, and 2 are mapped to priority group 0; dot1p priority 3 is mapped to priority group 1; dot1p priority 4 is mapped to priority group 2; dot1p priorities 5, 6, and 7 are mapped to priority group 4. All priorities that map to the same egress queue must be in the same priority group.	priority-pgid *dot1p0_group_num dot1p1_group_num dot1p2_group_num dot1p3_group_num dot1p4_group_num dot1p5_group_num dot1p6_group_num dot1p7_group_num*	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (**no priority pgid** command), the PFC and ETS parameters revert to their default values on the interfaces on which the DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.

  As a result, PFC and lossless port queues are disabled on 802.1p priorities, and all priorities are mapped to the same priority queue and equally share port bandwidth.
- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, first create a new DCB map with the desired PFC and ETS settings and apply the new map to the interfaces to override the previous DCB map settings. Then delete the original dot1p priority-priority group mapping.

  If you delete the dot1p priority-priority group mapping (**no priority pgid** command) before you apply the new DCB map, the default PFC and ETS parameters are applied on the interfaces. This change may create a DCB mismatch with peer DCB devices and interrupt network operation.

Applying a DCB map on server-facing Ethernet ports

You can apply a DCB map only on a physical Ethernet interface and can apply only one DCB map per interface.

Step	Task	Command	Command Mode
1	Enter interface configuration mode on a server-facing port or port channel to apply a DCB map. You cannot apply a DCB map on a port channel. However, you can apply a DCB map on the ports that are members of the port channel.	interface {**tengigabitEthernet** *slot/port* | **fortygigabitEthernet** *slot/port*}	CONFIGURATION
2	Apply the DCB map on an Ethernet port or port channel. The port is configured with the PFC and ETS settings in the DCB map; for example: FTOS# **interface tengigabitEthernet 0/0** FTOS(config-if-te-0/0)# **dcb-map SAN_DCB1** Repeat this step to apply a DCB map to more than one port or port channel.	dcb-map *name*	INTERFACE

Creating an FCoE VLAN

Create a dedicated VLAN to send and receive Fibre Channel traffic over FCoE links between servers and a fabric over an NPG. The NPG receives FCoE traffic and forwards de-capsulated FC frames over FC links to SAN switches in a specified fabric.

Step	Task	Command	Command Mode
1	Create the dedicated VLAN for FCoE traffic. Range: 2-4094. VLAN 1002 is commonly used to transmit FCoE traffic.	interface vlan *vlan-id*	CONFIGURATION

When you apply an FCoE map to an Ethernet port (Applying an FCoE map on server-facing Ethernet ports), the port is automatically configured as a tagged member of the FCoE VLAN.

Creating an FCoE map

An FCoE map consists of:

- An association between the dedicated VLAN used to carry FCoE traffic and SAN fabric where the storage arrays are installed. Use a separate FCoE VLAN for each fabric to which FCoE traffic is forwarded. Any non-FCoE traffic sent on a dedicated FCoE VLAN is dropped.
- The FC-MAP value used to generate the fabric-provided MAC address (FPMA). The FPMA is used by servers to transmit FCoE traffic to the fabric. You can associate an FC-MAP with only one FCoE VLAN and vice versa.
- FCF priority: the priority used by a server CNA to select an upstream FCoE forwarder (FCF)
- FIP keepalive (FKA) advertisement timeout

The values for the FCoE VLAN, fabric ID and FC-MAP must be unique. Apply an FCoE map on downstream server-facing Ethernet ports and upstream fabric-facing Fibre Channel ports.

Step	Task	Command	Command Mode
1	Create an FCoE map which contains parameters used in the communication between servers and a SAN fabric.	fcoe-map *map-name*	CONFIGURATION
2	Configure the association between the dedicated VLAN (configured in Creating an FCoE VLAN) and the fabric where the desired storage arrays are installed. The fabric and VLAN ID numbers must be the same. Fabric and VLAN ID range: 2-4094. For example: **fabric id 10 vlan 10**	fabric-id *fabric-num* vlan *vlan-id*	FCoE MAP
3	Add a text description of the settings in the FCoE map. Maximum: 32 characters.	description *text*	FCoE MAP
4	Specify the FC-MAP value used to generate a fabric-provided MAC address, which is required to send FCoE traffic from a server on the FCoE VLAN to the FC fabric specified in Step 2. You must enter a unique MAC address prefix as the FC-MAP value for each fabric. Range: 0EFC00–0EFCFF. Default: None.	fc-map *fc-map-value*	FCoE MAP
5	Configure the priority used by a server CNA to select the FCF for a fabric login (FLOGI). Range: 1-255. Default: 128.	fcf-priority *priority*	FCoE MAP
6	Enable the monitoring FIP keep-alive messages (if it is disabled) to detect if other FCoE devices are reachable. Default: FIP keep-alive monitoring is enabled.	keepalive	FCoE MAP
7	Configure the time interval (in seconds) used to transmit FIP keepalive advertisements. Range: 8-90 seconds. Default: 8 seconds.	fka-adv-period *seconds*	FCoE MAP

Applying an FCoE map on server-facing Ethernet ports

You can apply multiple FCoE maps on an Ethernet port or port channel. When you apply an FCoE map on a server-facing port or port channel:

- The port is configured to operate in hybrid mode (accept both tagged and untagged VLAN frames).
- The associated FCoE VLAN is enabled on the port or port channel.

When you enable a server-facing Ethernet port, the servers respond to the FIP advertisements by performing FLOGIs on upstream virtualized FCF ports. The NPG forwards the FLOGIs as fabric discovery (FDISC) messages to a SAN switch.

Step	Task	Command	Command Mode
1	Configure a server-facing Ethernet port or port channel with an FCoE map.	interface {**tengigabitEthernet** *slot*/*port* | **fortygigabitEthernet**	CONFIGURATION

Step	Task	Command	Command Mode
		slot/*port* | **port-channel** *num*}	
2	Apply the FCoE/FC configuration in an FCoE map on the Ethernet port. Repeat this step to apply an FCoE map to more than one port; for example: FTOS# **interface tengigabitEthernet 0/0** FTOS(config-if-te-0/0)# **fcoe-map SAN_FABRIC_A** FTOS# **interface port-channel 3** FTOS(config-if-po-3)# **fcoe-map SAN_FABRIC_A** FTOS# **interface fortygigabitEthernet 0/48** FTOS(config-if-fo-0/0)# **fcoe-map SAN_FABRIC_A**	fcoe-map *map-name*	INTERFACE or INTERFACE PORT_CHANNEL
3	Enable the port for FCoE transmission using the map settings.	no shutdown	INTERFACE

Applying an FCoE Map on fabric-facing FC ports

MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module FC ports are configured by default to operate in N port mode to connect to an F port on an FC switch in a fabric. You can apply only one FCoE map on an FC port.

When you apply an FCoE map on a fabric-facing FC port, the FC port becomes part of the FCoE fabric, whose settings in the FCoE map are configured on the port and exported to downstream server CNA ports.

Each MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module FC port is associated with an Ethernet MAC address (FCF MAC address). When you enable a fabric-facing FC port, the FCoE map applied to the port starts sending FIP multicast advertisements using the parameters in the FCoE map over server-facing Ethernet ports. A server sees the FC port with its applied FCoE map as an FCF port.

Step	Task	Command	Command Mode
1	Configure a fabric-facing FC port. For information on how to configure FC port settings, see Fibre Channel Interfaces.	interface fibrechannel *slot*/*port*	CONFIGURATION
2	Apply the FCoE and FC fabric configurations in an FCoE map to the port. Repeat this step to apply an FCoE map to more than one FC port; for example: FTOS# **interface fi 0/0** FTOS(config-if-fc-0/0)# **fabric SAN_FABRIC_A**	fabric *map-name*	INTERFACE FIBRE_CHANNEL
3	Enable the port for FC transmission.	no shutdown	INTERFACE FIBRE_CHANNEL

Tip: You can apply a DCB or FCoE map to a range of Ethernet or Fibre Channel interfaces by using the **interface range** command; for example:

```
FTOS(config)# interface range tengigabitEthernet 1/12 - 23 , tengigabitEthernet
2/24 - 35
```

```
FTOS(config)# interface range fibrechannel 0/0 - 3 , fibrechannel 0/8 - 11
```

Enter the keywords **interface range** followed by an interface type and port range. A port range must contain spaces before and after the dash. Separate each interface type and port range with a space, comma, and space as shown in the above examples.

Sample Configuration

1. Configure a DCB map with PFC and ETS settings:

```
FTOS(config)# dcb-map SAN_DCB_MAP
FTOS(config-dcbx-name)# priority-group 0 bandwidth 60 pfc off
FTOS(config-dcbx-name)# priority-group 1 bandwidth 20 pfc on
FTOS(config-dcbx-name)# priority-group 2 bandwidth 20 pfc on
FTOS(config-dcbx-name)# priority-group 4 strict-priority pfc off
FTOS(conf-dcbx-name)# priority-pgid 0 0 0 1 2 4 4 4
```

2. Apply the DCB map on a downstream (server-facing) Ethernet port:

```
FTOS(config)# interface tengigabitethernet 1/0
FTOS(config-if-te-0/0)#dcb-map SAN_DCB_MAP
```

3. Create the dedicated VLAN to be used for FCoE traffic:

```
FTOS(conf)#interface vlan 1002
```

4. Configure an FCoE map to be applied on downstream (server-facing) Ethernet and upstream (core-facing) FC ports:

```
FTOS(config)# fcoe-map SAN_FABRIC_A
FTOS(config-fcoe-name)# fabric-id 1002 vlan 1002
FTOS(config-fcoe-name)# description "SAN_FABRIC_A"
FTOS(config-fcoe-name)# fc-map 0efc00
FTOS(config-fcoe-name)# keepalive
FTOS(config-fcoe-name)# fcf-priority 128
FTOS(config-fcoe-name)# fka-adv-period 8
```

5. Enable an upstream FC port:

```
FTOS(config)# interface fibrechannel 0/0
FTOS(config-if-fc-0)# no shutdown
```

6. Enable a downstream Ethernet port:

```
FTOS(config)#interface tengigabitEthernet 0/0
FTOS(conf-if-te-0)# no shutdown
```

Displaying NPIV Proxy Gateway Information

To display information on NPG operation, use the show commands in the following table

Table 18. Displaying NPIV Proxy Gateway Information

Command	Description
show interfaces status	Displays the operational status of Ethernet and Fibre Channel interfaces on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG.
show fcoe-map [**brief** | *map-name*]	Displays the Fibre Channel and FCoE configuration parameters in FCoE maps. Enter the **brief** keyword to display an overview of currently configured FCoE maps. Enter the name of an FCoE map to display the FC and FCoE parameters configured in the map to be applied on MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet (FCoE) and FC ports.
show qos dcb-map *map-name*	Displays configuration parameters in a specified DCB map.
show npiv devices [**brief**]	Displays information on FCoE and FC devices currently logged in to the NPG.
show fc switch	Displays the FC mode of operation and worldwide node (WWN) name of an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module.

show interfaces status Command Example

```
FTOS# show interfaces status
Port     Description  Status Speed     Duplex Vlan
Fc 0/0                Up     8000 Mbit Auto   --
Fc 0/1                Down   Auto      Auto   --
Fc 0/2                Down   Auto      Auto   --
Fc 0/3                Down   Auto      Auto   --
Fc 0/4                Down   Auto      Auto   --
Fc 0/5                Down   Auto      Auto   --
Fc 0/6                Down   Auto      Auto   --
Fc 0/7                Down   Auto      Auto   --
Fc 0/8                Down   Auto      Auto   --
Fc 0/9                Down   Auto      Auto   --
Fc 0/10               Down   Auto      Auto   --
Fc 0/11               Down   Auto      Auto   --
Te 1/12               Down   Auto      Auto   --
Te 1/13               Down   Auto      Auto   --
Te 1/14               Down   Auto      Auto   --
Te 1/15               Down   Auto      Auto   --
Te 1/16               Down   Auto      Auto   --
Te 1/17               Down   Auto      Auto   --
Te 1/18               Down   Auto      Auto   --
Te 1/19               Up     10000 Mbit Full   --
Te 1/20               Down   Auto      Auto   --
Te 1/21               Down   Auto      Auto   --
```

Table 19. show interfaces status Field Descriptions

Field	Description
Port	Server-facing 10GbE Ethernet (Te), 40GbE Ethernet (Fo), or fabric-facing Fibre Channel (Fc) port with *slot*/*port* information.
Description	Text description of port.
Status	Operational status of port:

	Ethernet ports - up (transmitting FCoE and LAN storage traffic) or down (not transmitting traffic). Fibre Channel ports - up (link is up and transmitting FC traffic) or down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer SW port), or removed (port has been shut down).
Speed	Transmission speed (in Megabits per second) of Ethernet and FC iports, including auto-negotiated speed (Auto).
Duplex	Data transmission mode: Full (allows communication in both directions at the same time)Half (allows communication in both directions but not at the same time)Auto (auto-negotiated transmission).
VLAN	VLAN IDs of the VLANs in which the port is a member.

show fcoe-map Command Examples

```
FTOS# show fcoe-map brief
Fabric-Name   Fabric-Id         Vlan-Id  FC-MAP    FCF-Priority  Config-
State      Oper-State
fid_1003        1003            1003     0efc03    128
ACTIVE            UP
fid_1004        1004            1004     0efc04    128
ACTIVE            DOWN


FTOS# show fcoe-map fid_1003

Fabric Name         fid_1003
Fabric Id           1003
Vlan Id             1003
Vlan priority       3
FC-MAP              0efc03
FKA-ADV-Period      8
Fcf Priority        128
Config-State        ACTIVE
Oper-State          UP
Members
Fc 0/0
Te 0/14 Te 0/16
```

Table 20. show fcoe-map Field Descriptions

Field	Description
Fabric-Name	Name of a SAN fabric.
Fabric ID	The ID number of the SAN fabric to which FC traffic is forwarded.
VLAN ID	The dedicated VLAN used to transport FCoE storage traffic between servers and a fabric over the NPG. The configured VLAN ID must be the same as the fabric ID.
VLAN priority	FCoE traffic uses VLAN priority 3. (This setting is not user-configurable.)
FC-MAP	FCoE MAC-address prefix value - The unique 24-bit MAC address prefix that identifies a fabric.
FKA-ADV-period	Time interval (in seconds) used to transmit FIP keepalive advertisements.
FCF Priority	The priority used by a server to select an upstream FCoE forwarder.

Config-State	Indicates whether the configured FCoE and FC parameters in the FCoE map are valid: Active (all mandatory FCoE and FC parameters are correctly configured) or Incomplete (either the FC-MAP value, fabric ID, or VLAN ID are not correctly configured).
Oper-State	Operational status of link to the fabric: up (link is up and transmitting FC traffic), down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer FC port), or removed (port has been shut down).
Members	MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet and FC ports that are members of the dedicated FCoE VLAN that carries storage traffic to the specified fabric.

show qos dcb-map Command Examples

```
FTOS# show qos dcb-map dcbmap2

State   :Complete
PfcMode:ON
--------------------
PG:0 TSA:ETS  BW:50  PFC:OFF
Priorities:0 1 2 4 5 6 7


PG:1 TSA:ETS  BW:50  PFC:ON
Priorities:3
```

Table 21. show qos dcb-map Field Descriptions

Field	Description
State	Complete: All mandatory DCB parameters are correctly configured. In progress: The DCB map configuration is not complete. Some mandatory parameters are not configured.
PFC Mode	PFC configuration in the DCB map: On (enabled) or Off.
PG	Priority group configured in the DCB map.
TSA	Transmission scheduling algorithm used in the DCB map: Enhanced Transmission Selection (ETS).
BW	Percentage of bandwidth allocated to the priority group.
PFC	PFC setting for the priority group: On (enabled) or Off.
Priorities	802.1p priorities configured in the priority group.

show npiv devices brief Command Example

```
FTOS# show npiv devices brief

Total NPIV Devices = 2
-------------------------------------------------------------------------------
------------------------

ENode-Intf      ENode-WWPN           FCoE-Vlan  Fabric-Intf     Fabric-Map
LoginMethod     Status
-------------------------------------------------------------------------------
------------------------

Te 0/12      20:01:00:10:18:f1:94:20   1003      Fc 0/5
fid_1003           FLOGI     LOGGED_IN
```

```
Te 0/13      10:00:00:00:c9:d9:9c:cb   1003      Fc 0/0
fid_1003           FDISC     LOGGED_IN
```

Table 22. show npiv devices brief Field Descriptions

Field	Description
Total NPIV Devices	Number of downstream ENodes connected to a fabric over the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPIV proxy gateway.
ENode-Intf	MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet interface (*slot*/*port*) to which a server CNA is connected.
ENode-WWPN	Worldwide port name (WWPN) of a server CNA port.
FCoE-Vlan	VLAN ID of the dedicated VLAN used to transmit FCoE traffic to and from the fabric.
Fabric-Intf	Fabric-facing Fibre Channel port (*slot*/*port*) on which FC traffic is transmitted to the specified fabric.
Fabric-Map	Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.
Login Method	Method used by the server CNA to log in to the fabric; for example: FLOGI - ENode logged in using a fabric login (FLOGI). FDISC - ENode logged in using a fabric discovery (FDISC).
Status	Operational status of the link between a server CNA port and a SAN fabric: Logged In - Server has logged in to the fabric and is able to transmit FCoE traffic

show npiv devices Command Example

```
FTOS# show npiv devices
ENode[0]:
ENode MAC    :     00:10:18:f1:94:21
ENode Intf   :     Te 0/12
FCF MAC      :     5c:f9:dd:ef:10:c8
Fabric Intf  :     Fc 0/5
FCoE Vlan    :     1003
Fabric Map   :     fid_1003
ENode WWPN   :     20:01:00:10:18:f1:94:20
ENode WWNN   :     20:00:00:10:18:f1:94:21
FCoE MAC     :     0e:fc:03:01:02:01
FC-ID        :     01:02:01
LoginMethod  :     FLOGI
Secs         :     5593
Status       :     LOGGED_IN

ENode[1]:
ENode MAC    :     00:10:18:f1:94:22
ENode Intf   :     Te 0/13
FCF MAC      :     5c:f9:dd:ef:10:c9
Fabric Intf  :     Fc 0/0
FCoE Vlan    :     1003
Fabric Map   :     fid_1003
ENode WWPN   :     10:00:00:00:c9:d9:9c:cb
ENode WWNN   :     10:00:00:00:c9:d9:9c:cd
FCoE MAC     :     0e:fc:03:01:02:02
FC-ID        :     01:02:01
LoginMethod  :     FDISC
Secs         :     5593
```

```
Status        :      LOGGED_IN
```

Table 23. show npiv devices Field Descriptions

Field	Description
ENode [*number*]	Server CNA that has successfully logged in to a fabric over an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet port in ENode mode.
Enode MAC	MAC address of a server CNA port.
Enode Intf	Port number of a server-facing Ethernet port operating in ENode mode.
FCF MAC	Fibre Channel forwarder MAC: MAC address of MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module FCF interface.
Fabric Intf	Fabric-facing MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Fibre Channel port (*slot*/*port*) on which FCoE traffic is transmitted to the specified fabric.
FCoE VLAN	ID of the dedicated VLAN used to transmit FCoE traffic from a server CNA to a fabric and configured on both the server-facing MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module port and server CNA port.
Fabric Map	Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.
Enode WWPN	Worldwide port name of the server CNA port.
Enode WWNN	Worldwide node name of the server CNA.
FCoE MAC	Fabric-provided MAC address (FPMA). The FPMA consists of the FC-MAP value in the FCoE map and the FC-ID provided by the fabric after a successful FLOGI. In the FPMA, the most significant bytes are the FC-MAP; the least significant bytes are the FC-ID.
FC-ID	FC port ID provided by the fabric.
LoginMethod	Method used by the server CNA to log in to the fabric; for example, FLOGI or FDISC.
Secs	Number of seconds that the fabric connection is up.
State	Status of the fabric connection: logged in.

show fc switch Command Example

```
FTOS# show fc switch
Switch Mode : NPG
Switch WWN  : 10:00:5c:f9:dd:ef:10:c0
FTOS#
```

Table 24. show fc switch Command Description

Field	Description
Switch Mode	Fibre Channel mode o f operation of an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module. Default: NPG (configured as an NPIV proxy gateway).
Switch WWN	Factory-assigned worldwide node (WWN) name of the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module. The MXL

Field	Description
	10/40GbE Switch and M I/O Aggregator with the FC Flex IO module WWN name is not user-configurable.